arXiv:1601.03060v3 [q-bio.NC] 5 Dec 2016

Efficient Probabilistic Inference in Generic Neural Networks Trained with Non-Probabilistic Feedback

A. Emin Orhan[†] eorhan@cns.nyu.edu
Wei Ji Ma[†,‡] weijima@nyu.edu
†Center for Neural Science and ‡Department of Psychology
New York University, New York, NY 10003

December 6, 2016

Abstract

Animals perform near-optimal probabilistic inference in a wide range of psychophysical tasks. Probabilistic inference requires trial-to-trial representation of the uncertainties associated with task variables and subsequent use of this representation. Previous work has implemented such computation using neural networks with hand-crafted and task-dependent operations. We show that generic neural networks trained with non-probabilistic feedback using a simple error-based learning rule perform near-optimal probabilistic inference in nine common psychophysical tasks. In a probabilistic categorization task, error-based learning in a generic network simultaneously accounts for a monkey's learning curve and for the evolution of qualitative aspects of its choice behavior. In all tasks, the number of neurons required for a given level of performance grows sub-linearly with the size of the input population, a substantial improvement on previous implementations of probabilistic inference using hand-crafted neurons. The trained networks develop a novel sparsity-based probabilistic population code in a task-dependent manner. The tractability of the investigated tasks allows us to mechanistically understand the nature of this code. Our results suggest that probabilistic inference emerges naturally in generic neural networks trained with error-based learning rules and that these types of networks can be used as simple plausible neural models of probabilistic inference.

Introduction

When faced with noisy and incomplete sensory information, humans and other animals often behave near-optimally [2, 10, 21, 23, 34, 46]. Optimal behavior requires that the brain compute posterior distributions over task-relevant variables, which often involves complex operations such as multiplying probability distributions or marginalizing over latent variables. How do neural circuits implement such operations? A prominent framework addressing this question is the probabilistic population coding (PPC) framework, according to which the population activity on a single trial encodes a probability distribution rather than a single estimate and computations with probability distributions can be carried out by suitable operations on the corresponding neural responses [28, 48]. For example, Ma et al. [28] showed that if neural variability belongs to a particular class of probability distributions, the posterior distribution in cue combination tasks can be computed with a linear combination of the input responses. Moreover, in this scheme, the form of neural variability is preserved between the input and the output, leading to an elegantly modular code. In more complex tasks, linear operations are insufficient and it has been argued that multiplication and division of neural responses are necessary for optimal inference [3, 4, 29, 30, 38].

Upon closer look, however, these previous implementations of PPC suffer from several shortcomings. First, the networks in these studies were either fully manually designed, or partially manually designed and partially trained with large amounts of probabilistic data to optimize explicitly probabilistic objectives, e.g. minimization of Kullback-Leibler (KL) divergence. Therefore, this literature does not address the important question of learning: how can probabilistic inference be learned from a realistic amount and type of data with minimal manual design of the networks? Second, although there are some commonalities between the neural operations required to implement probabilistic inference in different tasks, these operations generally differ from task to task. For instance, it has been argued that some form of divisive normalization of neural responses is necessary in tasks that involve marginalization [4]. However, the specific form of divisive normalization that individual neurons have to perform differs substantially from task to task. Therefore, it is unclear if probabilistic inference can be implemented in *generic* neural networks, whose neurons all perform the same type of neurally plausible operation. Third, in these studies, the number of neurons used for performing probabilistic inference scales unfavorably with the size of the input population (linearly in the case of cue combination, but at least quadratically in all other tasks). Therefore, the question of whether these tasks can be implemented more efficiently remains open.

In this paper, we address these issues. We show that generic neural networks trained with non-probabilistic error-based feedback perform near-optimal probabilistic inference in tasks with both categorical and continuous outputs. Generic neural networks of the type we use in this paper have a long history [39, 45, 49], and have recently been linked directly to cortical responses [7, 33, 43, 47]. Our main contribution is to connect generic neural networks to near-optimal probabilistic inference in common psychophysical tasks. For these tasks, we analyze the network generalization performance, the efficiency of the networks in terms of the number of neurons needed to a achieve a given level of performance, the nature of the emergent probabilistic population code and the mechanistic insights that can be gleaned from the trained networks. We also investigate whether the time-course of error-based learning in generic neural networks is realistic for non-linguistic animals learning to perform a probabilistic inference task.

Results

Tasks. We trained generic feedforward or recurrent neural networks on nine probabilistic psychophysical tasks that are commonly studied in the experimental and computational literature. The main tasks were: (a) linear cue combination [2, 10, 21], (b) coordinate transformations [4, 49], (c) Kalman filtering [25, 34, 46], (d) causal inference [23,30], (e) stimulus demixing [4], (f) binary categorization [38], (g) visual search with heterogeneous distractors [29] (see *Methods* for task details). We also trained generic networks on two additional "modular" tasks to be discussed below.

Networks. The networks all received noisy sensory information about the stimulus or the stimuli in the form of a neural population with Poisson variability. The hidden units of the networks were modeled as rectified linear units (ReLUs). ReLUs are commonly used in neural networks due to their demonstrated advantage over alternative non-linearities in gradient-based learning algorithms [13]. Linear (sigmoidal) output units were used in tasks with continuous (categorical) outputs. Schematic diagrams of the networks used for the main tasks considered in this paper are shown in Figure 1. Differences between the network architectures are due entirely to differences in the

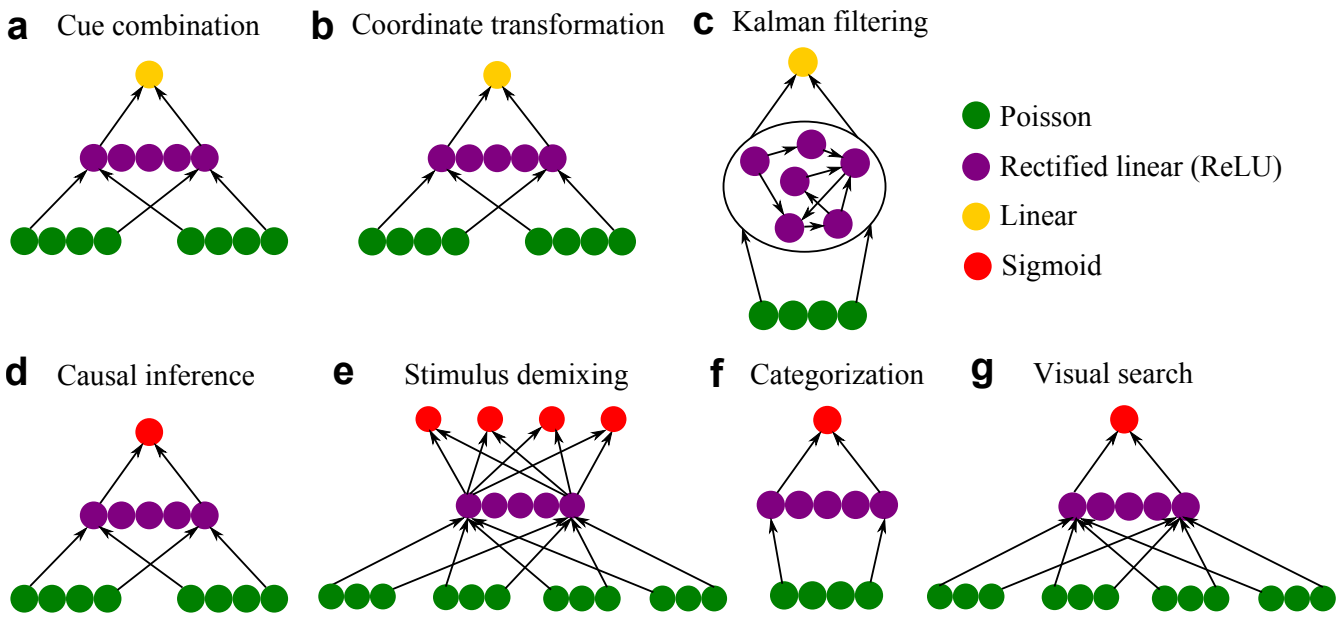

Figure 1: Network architectures used for the main tasks considered in this paper. Different colors represent different types of units. For tasks with continuous output variables (a-c), linear output units were used; whereas for tasks with categorical output variables (d-g), sigmoidal output units were used.

input and output requirements of different tasks: different tasks have different numbers of inputs or outputs and the outputs are continuous or categorical in different tasks. Other than these task-dictated differences, the networks are truly generic in the sense that they are composed of neurons that perform the same type of biologically plausible operations in all tasks.

Networks were trained to minimize mean squared error or cross-entropy in tasks with continuous or categorical outputs, respectively. Importantly, the networks were provided only with the actual stimulus values or the correct class labels as feedback in each trial. Thus, they did not receive any explicitly probabilistic feedback, nor were they explicitly trained to perform probabilistic inference.

We manipulated sensory reliability trial-by-trial via gain variables g multiplying the responses of the input populations, with higher gains corresponding to more reliable sensory information. In each task, networks were tested with a wide range of gains or gain combinations (in tasks with more than a single stimulus). To test the generalization capacity of the networks, we trained them with a limited range of gains or gain combinations, as well as with the full range of test gains or gain combinations. The latter unrestricted training regime is called the "All $\mathbf{g}$" condition, whereas the former limited training regime is called the "Restricted $\mathbf{g}$" condition in what follows. The specific gain ranges and gain combinations used in each task are indicated below (*Methods*).

Generic neural networks trained with error-based learning rules implement probabilistic inference in standard psychophysical tasks. The performance of well-trained networks is shown in Figure 2c and d for both "All $\mathbf{g}$" (black) and "Restricted $\mathbf{g}$" (red) training conditions in all tasks (complete learning curves of the networks are shown in supplementary Figure S1). For continuous tasks, performance is measured in terms of fractional RMSE defined as $100 \times (\mathrm{RMSE}_{\mathrm{netw}} - \mathrm{RMSE}_{\mathrm{opt}})/\mathrm{RMSE}_{\mathrm{opt}}$ where $\mathrm{RMSE}_{\mathrm{netw}}$ is the root mean-squared error (RMSE) of the trained network, $\mathrm{RMSE}_{\mathrm{opt}}$ is the RMSE of the posterior mean estimate. For categorical tasks, performance is measured in terms of fractional information loss defined as the average KL-divergence between the actual posterior and the network's output normalized by the mutual information between the class labels and the neural responses [38]. With this measure, a network that exactly reproduces the posterior achieves 0% information loss, whereas a network that produces random responses according to the prior probabilities of the classes has 100% information loss.

To make sure that optimal performance in our tasks cannot be easily mimicked by heuristic, non-probabilistic models, we also calculated the performance of non-probabilistic reference models that did not take the reliabilities of the inputs into account (Figure 2c-d, blue). The large performance gaps between these non-probabilistic models and

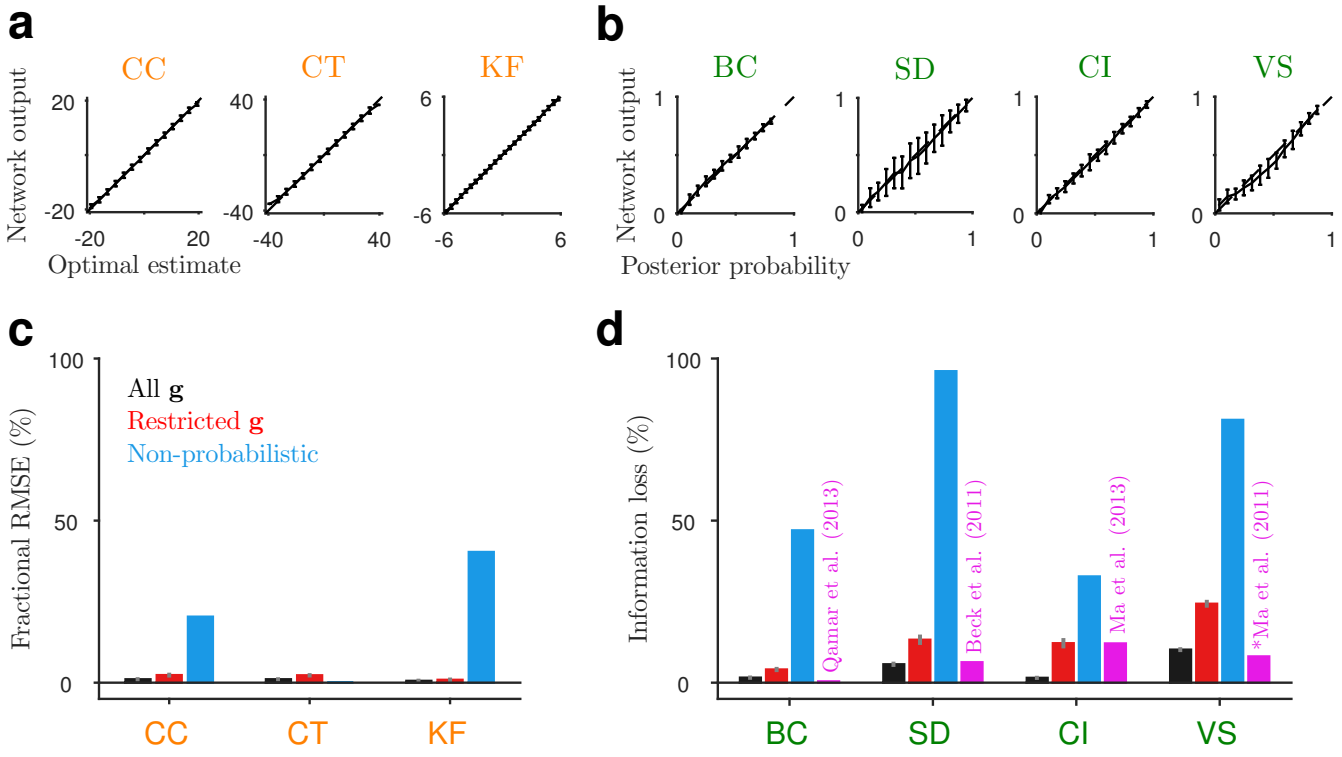

Figure 2: Performance of well-trained networks in the main tasks. (a) Optimal estimates vs. the network outputs in "all **g**" conditions of the continuous tasks. Error bars represent standard deviations. (b) Posterior probability of a given choice vs. the network output in "all **g**" conditions of the categorical tasks. Error bars represent standard deviations. For the stimulus demixing task (SD), where the output was four-dimensional, only the marginal posterior along the first dimension is shown for clarity. (c) Performance in continuous tasks. (d) Performance in categorical tasks. Blue bars show the performance of non-probabilistic heuristic models that do not take uncertainty into account. Note that optimal performance in the coordinate transformation task (CT) does not require taking uncertainty into account (see *Methods*). Magenta bars show the performance of hand-crafted networks in categorical tasks reported in earlier works. The asterisk in the visual search task (VS) indicates that the information loss value reported in [29] should be taken as a lower bound on the actual information loss, since they were not able to build a single network that solved the full visual search task in that paper. In (c) and (d), error bars (gray) represent means and standard errors over 10 independent runs of the simulations. Abbreviations: CC (cue combination); CT (coordinate transformation); KF (Kalman filtering); BC (binary categorization); SD (stimulus demixing); CI (causal inference); VS (visual search). Categorical tasks are labeled in green, continuous tasks in orange.

the optimal model suggest that approaching optimal performance in our tasks requires performing truly probabilistic inference, i.e. taking the reliabilities of the inputs into account.

In categorical tasks, the output nodes of the networks approximate the posterior probabilities of the classes given the inputs (Figure 2b, d). Theoretical guarantees ensure that this property holds under general conditions with a wide range of loss functions [18] (*Discussion*).

Encoding of posterior width in the hidden layers. For continuous tasks, training with mean squared error loss guarantees asymptotic convergence to the posterior mean. Do the networks also represent information about the posterior uncertainty in their hidden layers or do they discard this information? Representation of posterior uncertainty is evident in the Kalman filtering task where accurate encoding of the posterior mean at a particular moment already requires the encoding of the posterior mean and the posterior width at the previous moment and the optimal integration of these with the current sensory information in the recurrent activity of the network.

For the linear cue combination and coordinate transformation tasks, to test for the representation of posterior uncertainty in the hidden layer, we plugged the trained hidden layers into a network incorporating an additional

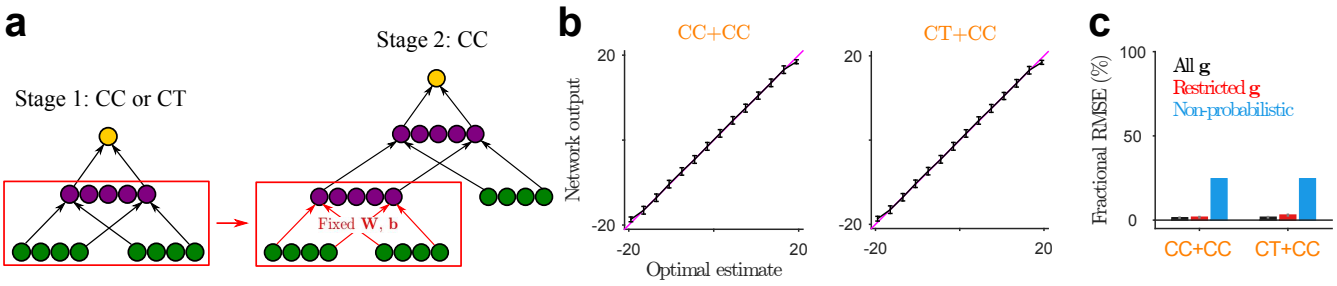

Figure 3: (a) Schematic diagram illustrating the design of the modular tasks. In stage 1, a single hidden layer network is trained on either cue combination (CC) or coordinate transformation (CT). In stage 2, the hidden layer of the trained network is plugged into another network. The combined network is trained on a three-input cue combination task with the parameters of the first network fixed. (b) Optimal estimates vs. network outputs for the three-input cue combination task. Error bars represent standard deviations. CT+CC is the case where the first network is trained on the coordinate transformation task, CC+CC is the case where the first network is trained on the two-input cue combination task. (c) Fractional RMSEs. Error bars represent standard errors over 10 independent runs of the simulations.

input population and fixed their parameters (Figure 3a). The rest of the network was then trained on a linear cue combination task with three input populations (a similar modular task was designed in [31]). If the fixed hidden layers do not encode the posterior width for the first two inputs, the combined network cannot perform the three-input cue combination task optimally. However, the combined networks were able to perform the three-input cue combination task with little information loss despite receiving information about the first two inputs only through the fixed hidden layers (Figure 3b-c). This suggests that although the initial networks were trained to minimize mean squared error, and hence were asymptotically guaranteed to reproduce the posterior means only, information about the posterior widths was, to a large extent, preserved in the hidden layer. The precise format in which posterior uncertainty is represented in the hidden layer activity will be discussed below. The combined coordinate transformation-cue combination (CT+CC) network also illustrates the generic nature of the representations learned by the hidden layers of our networks: the hidden layer of a network trained on the coordinate transformation task can be combined, without modification, with an additional input population to perform a different task, i.e. cue combination in this example.

Generalization to untrained stimulus conditions. Truly Bayesian computation requires that the components of the Bayesian computation, i.e. sensory likelihoods and the prior, be individually meaningful to the brain [32]. Thus, if we replace a particular likelihood for another, the system should continue to perform near-optimally. We tested for a limited form of such "Bayesian transfer" by examining whether the trained networks generalize to unseen values or combinations of sensory reliability ("restricted $\mathbf{g}$" conditions). As shown in Figure 2c and d (red bars), the networks were able to generalize well beyond the training conditions in all tasks. An example is shown in Figure 4a for the cue combination task. In this example, we trained a network with only two gain combinations, $\mathbf{g} \equiv (g_1, g_2) = (5, 5)$ and $\mathbf{g} = (25, 25)$, and tested it on all gain combinations of the form (g_1, g_2) where $g_1, g_2 \in \{5, 10, 15, 20, 25\}$ with up to five-fold gain differences between the two input populations (note that these gains are higher than those used in the main simulations to make the optimal combination rule approximately linear). To demonstrate that the trained networks performed qualitatively correct probabilistic inference, we set up cue conflict conditions similar to the cue conflict conditions in psychophysical studies [10], where we presented slightly different stimuli to the two input populations and manipulated the degree of conflict between the cues. The weights assigned to the first cue as a function of the gain ratio, g_1/g_2, are shown in Figure 4a both for the network and for the optimal rule. The network achieved low generalization error (fractional RMSE: 10.9%) even after as few as 50 training examples in the impoverished training condition and performed qualitatively correct probabilistic inference in the untrained conditions. In particular, the network correctly adjusted the weights assigned to the two cues even as the ratio of their reliabilities varied over a 25-fold range (Figure 4a).

The successful generalization performance of the neural networks is a result of two factors. First, the target function is approximately invariant to gain manipulations that differ between the training and test conditions. In cue combination, for instance, the target function is approximately invariant to the scaling of the input populations by arbitrary gains g_1 and g_2 (Equation 5). The second factor is the network's inductive biases, i.e. how it tends to behave outside the training domain. These inductive biases depend on the details of the network architecture [37].

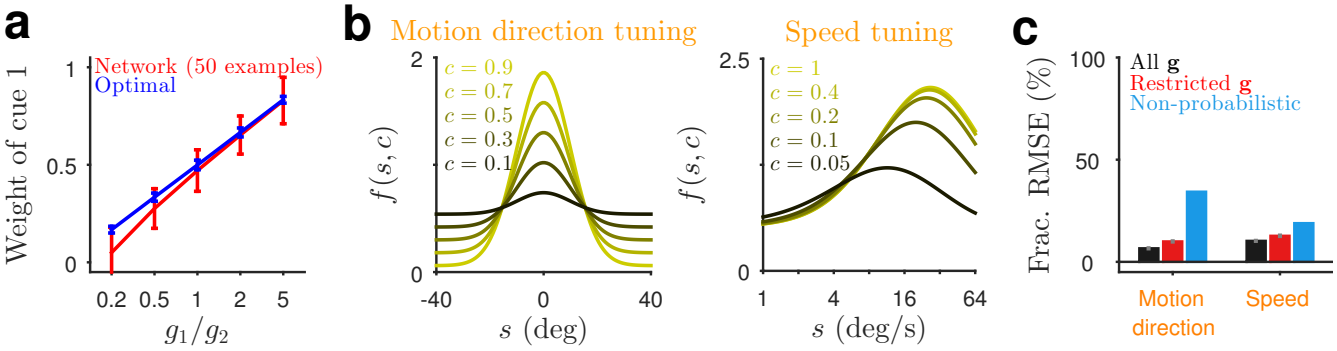

Figure 4: (a) Generalization capacity of neural networks in the cue combination task. The weights assigned to cue 1 in the cue conflict conditions as a function of the ratio of the input gains, g_1/g_2. In cue conflict trials, s_1 was first drawn randomly between -4 and 4, then s_2 was generated as $s_1+\Delta$ where Δ was one of $-16, -12, -8, -4, 4, 8, 12, 16$. The cue weight assigned by the network was calculated through the equation: $\hat{s} = w\hat{s}_{1,\text{opt}} + (1-w)\hat{s}_{2,\text{opt}}$, where $\hat{s}$ is the network output, $\hat{s}_{1,\text{opt}}$ and $\hat{s}_{2,\text{opt}}$ are the optimal estimates of s_1 and s_2. Note that we use relatively high gains in this particular example to make the optimal combination rule approximately linear. In the low gain conditions used in the main simulations, the optimal combination rule is no longer linear. The network experienced only 50 training examples in the restricted **g** condition with non-conflicting cues, hence it did not see any of the conditions shown here during training. Error bars represent standard deviations over 1000 simulated trials. (b) Tuning functions for motion direction as reported in [11, 35] and for speed as reported in [24], where the stimulus contrast or coherence variable c does not act multiplicatively on the mean firing rates. (c) Fractional RMSEs in cue combination tasks with tuning functions shown in (b). The networks perform near-optimal probabilistic inference regardless of the specific form in which stimulus reliability is encoded in the input populations.

Alternative representations of stimulus reliability in the input populations. Thus far, we have assumed that sensory reliability has a purely multiplicative effect on the responses of input neurons. Although this assumption likely holds for the effect of contrast on orientation selectivity in visual cortex [40], it is known to be violated for the effect of motion coherence on direction selectivity [11, 35] and for the effect of contrast on speed selectivity [24], and is unlikely to hold in the general case. The importance of this observation is that the Poisson-like PPC approach proposed in [28] cannot easily handle cases where "nuisance variables" such as contrast or coherence do not have a purely multiplicative effect on neural responses. By contrast, our approach does not make any assumptions about the representation of stimulus reliability in the input populations. We demonstrated this in two cases (Figure 4b-c): (i) cue combination with tuning functions of the form reported in [11, 35] where stimulus coherence affects both the gain and the baseline of the responses (Figure 4b, left) and (ii) cue combination with tuning functions of the form reported in [24] for speed where both the peak response and the preferred speed depend on stimulus contrast (Figure 4b, right). Thus, our networks can perform near-optimal inference regardless of the specific form in which stimulus reliability is encoded in the input populations.

Sparsity-based representation of posterior uncertainty. We now discuss how posterior uncertainty is represented in the hidden layers of the trained networks. We first note that in tasks with continuous output variables, the optimal solution is invariant to a multiplicative scaling g of the input responses (*Methods*, Equations 5-8). In such gain-invariant (or approximately gain-invariant) tasks, posterior uncertainty is represented in the sparsity of hidden layer activity. To understand the mechanism through which this sparsity-based representation arises, we investigated the conditions under which the network's output would be invariant to input gain scalings. We first derived an approximate analytical expression for the mean response of a typical hidden unit $\bar{\mu}$, as a function of the input gain g, the mean input μ to the hidden unit for unit gain, and the mean μ_b and the standard deviation σ_b of the biases of the hidden units (*Methods*). To minimize the dependence of the mean hidden unit response on g, we introduced the following measure of the total sensitivity of $\bar{\mu}$ to variations in g:

$$T_{\text{var}} = \int_{g_{\min}}^{g_{\max}} |\bar{\mu}'(g)| dg$$

where the prime represents the derivative with respect to g, and numerically minimized T_{var} with respect to μ, μ_b and σ_b, subject to the constraint that the mean response across different gains be equal to a positive constant K.

T_{var} was minimized for a negative mean input μ, positive μ_b and a large σ_b value (black star in Figure 5a). As an approximate rule, decreasing μ and increasing μ_b or σ_b lead to smaller T_{var} values. Because the input responses are always non-negative, the only way μ can be negative in our networks is if the mean input-to-hidden layer weight is negative. The negativity of the mean input μ implies that as the gain g increases, the distribution of the total input to the unit shifts to the left (Figure 5b, top), causing a larger proportion of the distribution to remain below the threshold (represented by the dashed line in Figure 5b), hence decreasing the probability that the neuron will have a non-zero response (Figure 5c, top). This mechanism causes the sparsification of the hidden unit responses with increasing g and is particularly effective when the mean bias, μ_b, is large and positive. Because increasing g also increases the variance of the total input to the unit, the mean response for those inputs that do cross the threshold increases (Figure 5d, top). As a result, the mean response of the neuron, which is a product of these two terms, remains roughly constant (Figure 5e, top).

We demonstrate this sparsification mechanism for a network trained on the coordinate transformation task in Figure 5f-i. Because the coordinate transformation task is approximately gain-invariant (*Methods*, Equation 6), the input-to-hidden layer weight distribution in the trained network was skewed toward negative values (Figure 5f) and the mean bias of the hidden units, μ_b, was positive (Figure 5g), as predicted from our simple mean-field model. Consequently, we found a strong positive correlation between the sparsity of hidden layer responses and the mean input response ($r = 0.81, P < 10^{-6}$; Figure 5i), but no correlation between the mean hidden layer response and the mean input response ($r = 0.09, P > 0.05$; Figure 5h).

The same type of analysis applies to the categorical tasks as well. However, the difference is that for some of our tasks with categorical outputs, in the optimal solution, the net input to the output unit had a strong dependence on g. For example, in causal inference, the input to the sigmoidal output unit scales approximately linearly with g (Equation 10). Similarly, in visual search, both global and local log-likelihood ratios have a strong dependence on g (through $\mathbf{r}_i$ in Equations 16-17). We emphasize that the distinction between g-dependence and g-invariance is not categorical: different tasks can have varying degrees of g-invariance or g-dependence and parameter choices in the same task can affect its g-dependence.

In the bottom panel of Figure 5b-e, predictions from the mean-field model are shown for a parameter combination where both μ and μ_b are small and slightly negative (represented by the magenta dot in Figure 5a). This parameter combination roughly characterizes the trained networks in the causal inference task (Figure 5j-m). In this case, because both μ and μ_b are close to 0, the probability of non-zero responses as a function of g stays roughly constant (Figure 5c, bottom), causing the mean response to increase with g (Figure 5e, bottom).

Based on our simple mean-field model, we therefore predicted that for those tasks where the net input to the output unit is approximately g-invariant, there should be a positive correlation between the sparsity of hidden unit responses and the input gain and no (or only a weak) correlation between the mean hidden unit response and the input gain. On the other hand, in tasks such as causal inference, where the net input to the output unit has a strong g-dependence, we predicted a positive correlation between the mean hidden unit response and the input gain and no (or only a weak) correlation between the sparsity of hidden unit responses and the input gain. We tested these predictions on our trained networks and confirmed that they were indeed correct (Figure 6a-b). For causal inference, visual search and stimulus demixing tasks, the correlation between the mean input and the sparsity of hidden layer responses was weak (Figure 6f), whereas for the remaining tasks, it was strong and positive. The opposite pattern was seen for the correlation between the mean input and the mean hidden layer response (Figure 6a). In g-dependent tasks such as causal inference, posterior uncertainty is thus represented largely in the mean hidden layer activity; whereas in approximately g-invariant tasks such as coordinate transformation, it is represented largely in the sparsity of hidden layer activity. The sparsity-based representation of posterior uncertainty in g-invariant tasks was again driven by large negative mean input-to-hidden layer weights and large positive mean biases (Figure 6c-d).

The difference between these two types of tasks (g-invariant and g-dependent) was also reflected in the tuning functions that developed in the hidden layer of the networks. For approximately g-invariant tasks such as coordinate transformation, increasing the input gain g sharpens the tuning of the hidden units (Figure 7a), whereas for g-dependent tasks such as causal inference, input gain acts more like a multiplicative factor scaling the tuning functions without changing their shape (Figure 7b).

We finally note that these results are decoder dependent. In the continuous tasks, for example, if we use a divisively normalized decoder instead of a linear read-out, posterior uncertainty is encoded in the mean hidden layer response, rather than in the sparsity of hidden layer responses (supplementary Figure S4).

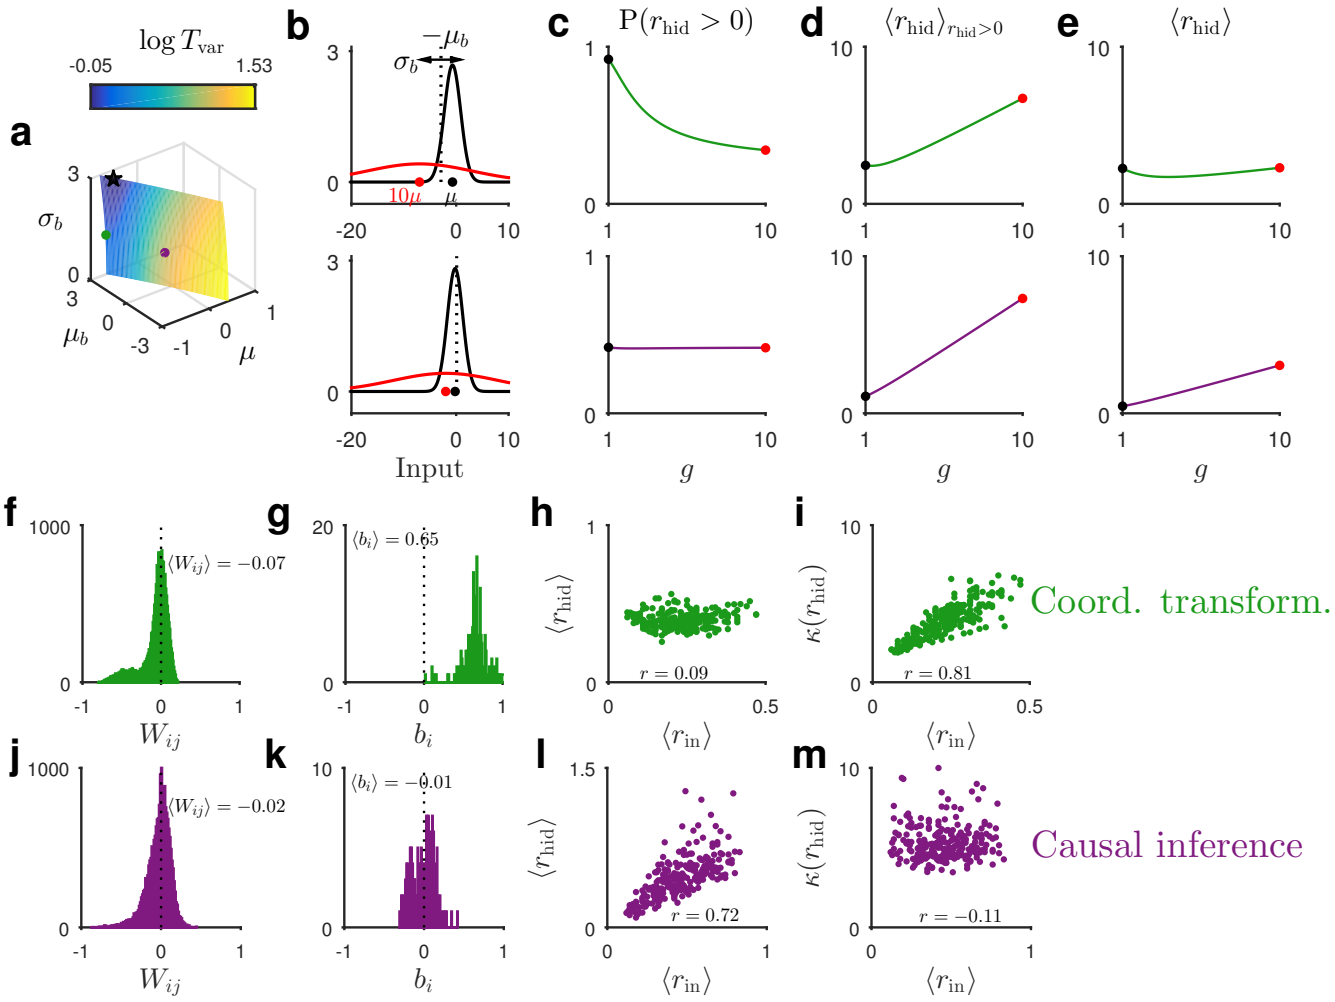

Figure 5: The mechanism underlying the sparsity-based representation of posterior uncertainty. (a) The variability index T_{var} (plotted in log units) as a function of the parameters μ, μ_b and σ_b for the constraint surface corresponding to $K = 2$. The optimal parameter values within the shown range are represented by the black star. The green and magenta dots roughly correspond to the parameter statistics in the trained coordinate transformation and causal inference networks, respectively. (b) For the parameter combination corresponding to the green dot, the mean input μ is negative and the mean bias μ_b is positive. Increasing the gain g thus shifts the input distribution to the left: compare the black and red lines for input distributions with a small and a large gain, respectively. The means of the input distributions are indicated by small dots underneath. (c) This, in turn, decreases the probability of non-zero responses, but (d) increases the mean of the non-zero responses; hence (e) the mean response of the hidden units, being a product of the two, stays approximately constant as the gain is varied. In the bottom panel of (b-e), the results are also shown for a different parameter combination represented by the magenta dot in (a). This parameter combination roughly characterizes the trained networks in the causal inference task. (f) For a network trained in the coordinate transformation task, distributions of the input-to-hidden layer weights, W_{ij}; (g) the biases of the hidden units, b_i; (h) scatter plot of the mean input $\langle r_{\text{in}} \rangle$ vs. the mean hidden unit activity $\langle r_{\text{hid}} \rangle$; (i) the scatter plot of the mean input vs. the kurtosis of hidden unit activity $\kappa(r_{\text{hid}})$. (j-m) Similar to (f-i), but for a network trained in the causal inference task.

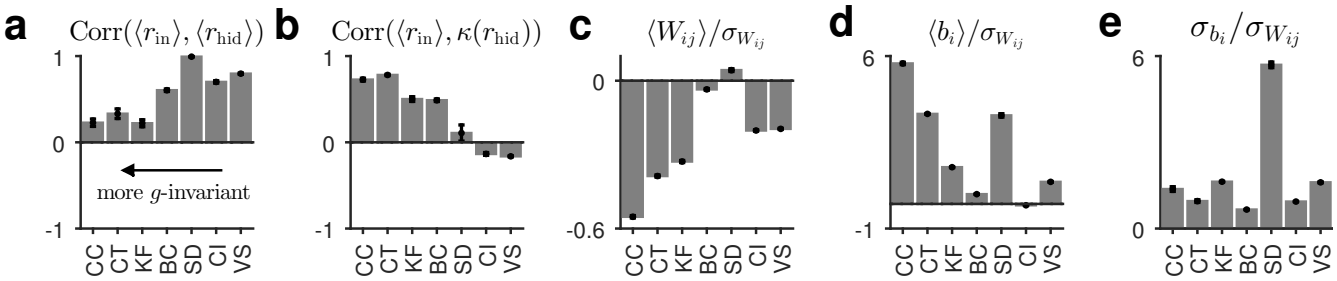

Figure 6: Encoding of posterior uncertainty and parameter statistics in the trained networks. (a) Trial-by-trial correlation between mean hidden unit response and mean input response; (b) trial-by-trial correlation between the sparsity (kurtosis) of hidden layer activity and mean input response; (c) mean input-to-hidden unit weight; (d) mean bias of the hidden units; (e) standard deviation of hidden unit biases. Parameter statistics are reported in units of the standard deviation of the input-to-hidden layer weights to make them consistent with the mean field analysis, where all variables are measured in units of the standard deviation of the input. CC: cue combination; CT: coordinate transformation; KF: Kalman filtering; BC: binary categorization; SD: stimulus demixing; CI: causal inference; VS: visual search. Error bars represent means and standard errors over 10 independent runs of the simulations.

Random networks. To investigate the architectural constraints on the networks capable of performing near-optimal probabilistic inference, we considered an alternative architecture, in which the input-to-hidden layer weights and the biases of the hidden units were set randomly and left untrained; only the hidden-to-output layer weights and the biases of the output units were trained (Figure 8a). Such random networks can be plausible models of some neural systems [8, 42]. Given the same amount of computational resources, these networks performed substantially worse than the fully trained networks (Figure 8b). A well-known theoretical result can explain the inefficiency of random networks [1]: the approximation error of neural networks with adjustable hidden units scales as $O(1/n)$ with n denoting the number of hidden units, whereas for networks with fixed hidden units, as in our random networks, the scaling is much worse: $O(1/n^{2/d})$ where d is the dimensionality of the problem, suggesting that they need exponentially more neurons than fully trained networks in order to achieve the same level of performance.

Making the networks biologically more plausible: recurrent EI networks. So far, we have only considered feedforward networks with undifferentiated neurons. To investigate whether introducing more biological realism would severely constrain the capacity of the networks to perform near-optimal probabilistic inference, following the approach proposed in [41], we trained fully recurrent networks with separate excitatory and inhibitory populations that respected Dale's law: excitatory neurons projecting only positive weights, inhibitory neurons only negative weights. The input populations were also divided into excitatory and inhibitory sub-populations that obeyed Dale's law (Figure 8c). The performance of these recurrent EI networks were slightly worse than, but not substantially different from, the corresponding fully-trained feedforward networks (Figure 8d). Moreover, the networks re-coded stimulus reliability in a similar manner to the feedforward networks (supplementary Figure S2), suggesting that the main results reported for feedforward networks are robust to the incorporation of more biological realism into our networks.

Error-based learning in generic neural networks accounts for the evolution of behavioral choice and accuracy in a binary categorization task. The dependence of the networks' performance on the number of training trials (Figure S1) suggests a possible explanation for deviations from optimal inference sometimes observed in experimental studies: i.e. insufficient training in the task. Testing this hypothesis rigorously is complicated by possible prior exposure of the subjects to similar stimuli or tasks under natural conditions. Among the tasks considered in this paper, the binary categorization task minimizes such concerns, because it involves classifying stimuli into arbitrary categories. Moreover, in this task, the behavior of both human and monkey observers were best explained by heuristic models that were quantitatively suboptimal, but qualitatively consistent with the optimal inference model [38]. Therefore, we sought to test the insufficient-training hypothesis for suboptimal inference in this task.

The stimulus distributions for the two categories and the decision boundaries predicted by the optimal (OPT) and three suboptimal models (FIX, LIN, QUAD) are shown in Figure 9a-b. Different suboptimal models make different assumptions about the dependence of the decision boundary on the sensory noise, σ (Figure 9b). In particular,

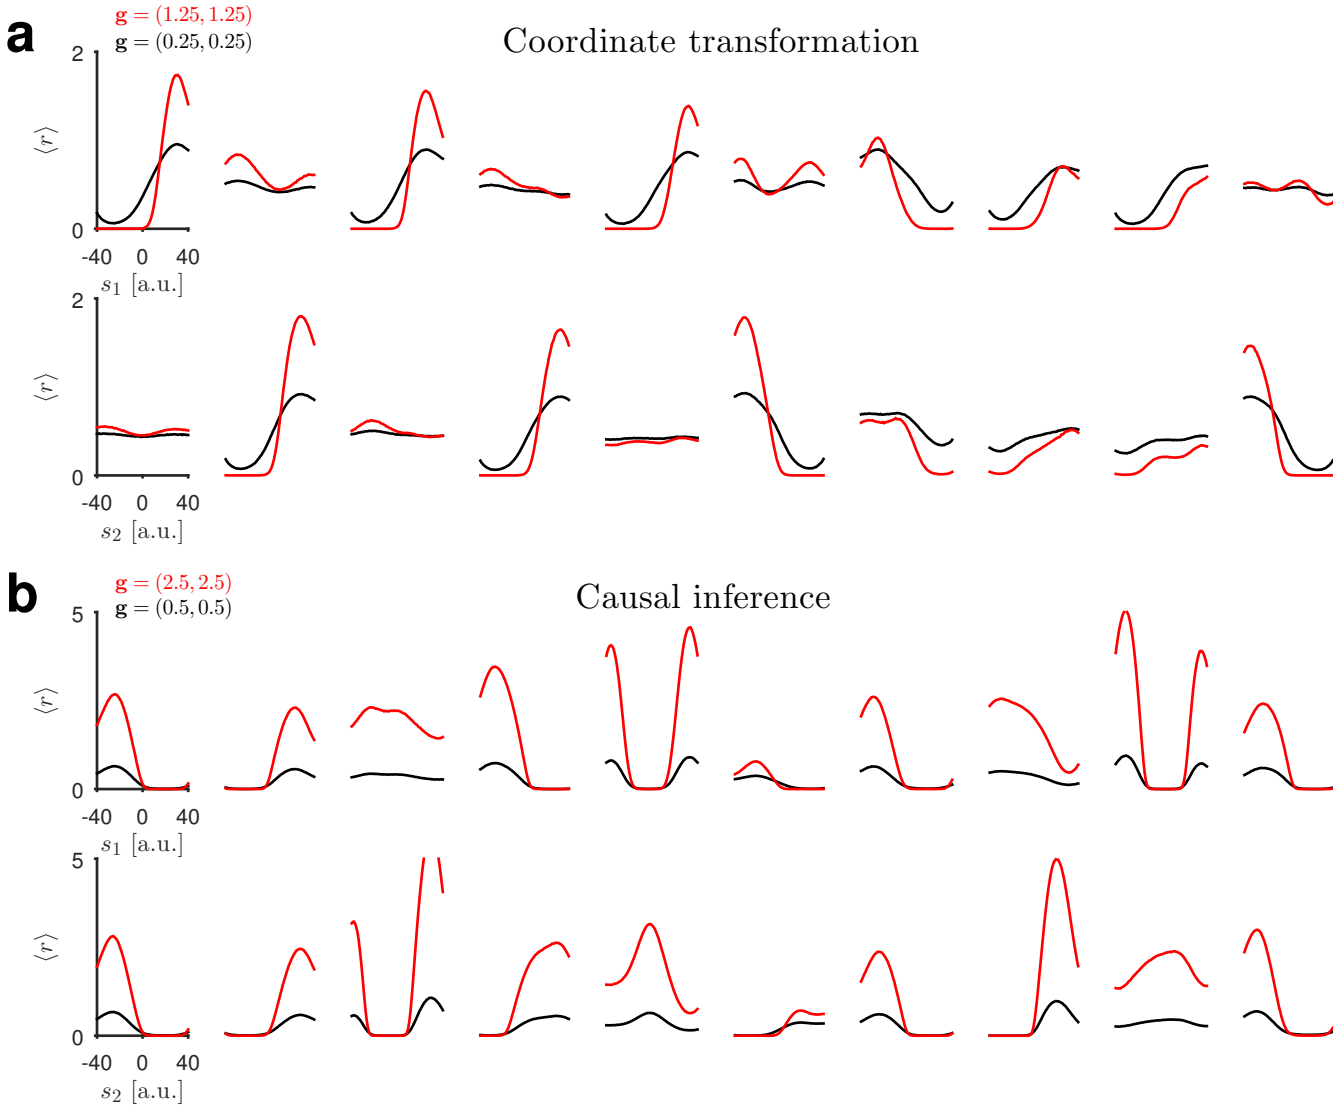

Figure 7: One-dimensional tuning functions of 10 representative hidden units at two input gains (a) in a network trained in the coordinate transformation task and (b) in a network trained in the causal inference task. In both (a) and (b), the first row shows the tuning functions with respect to s_1 (averaged over s_2) and the second row shows the tuning functions with respect to s_2. Increasing the gain sharpens the tuning curves in the coordinate transformation task, whereas it has a more multiplicative effect in the causal inference task.

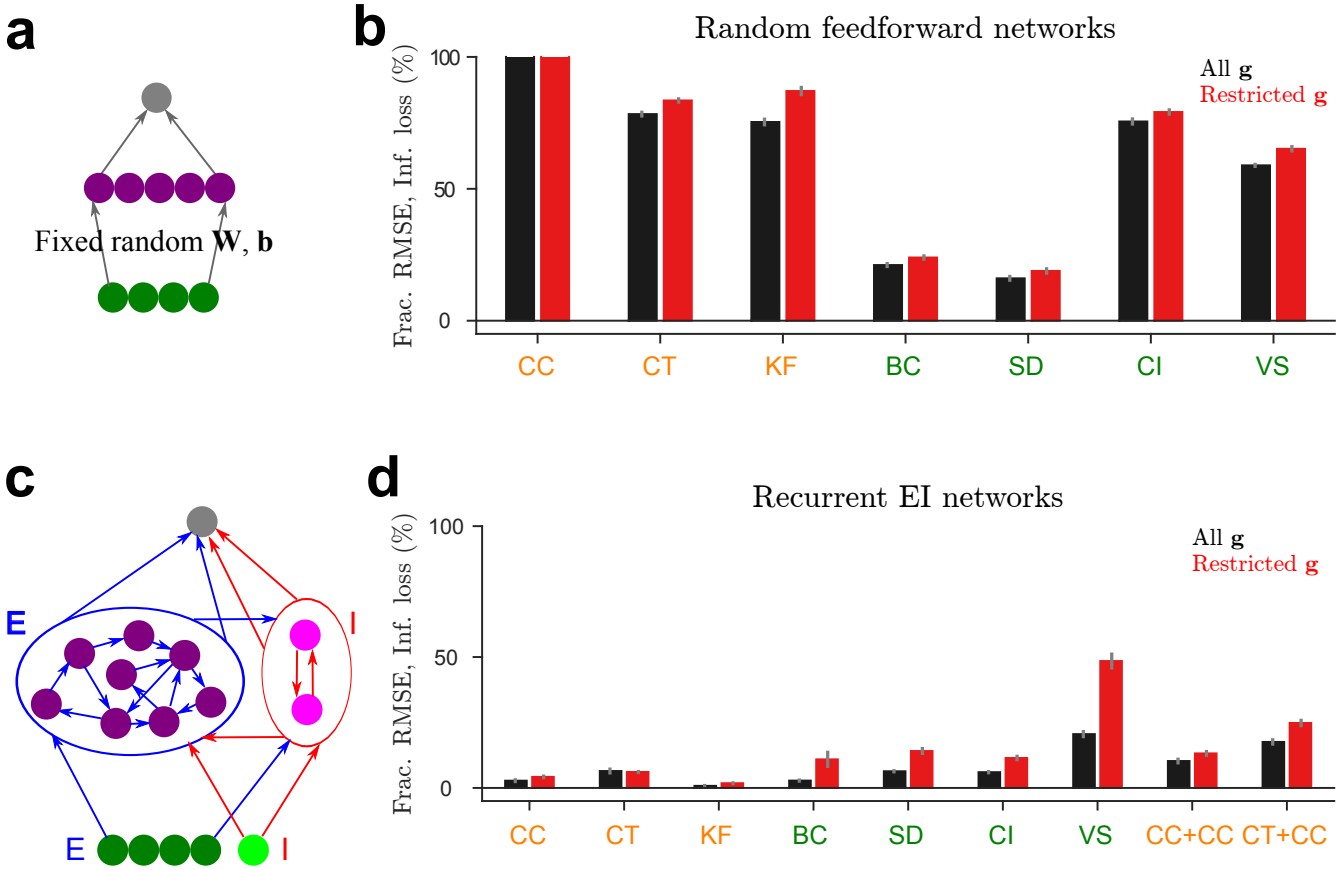

Figure 8: Architectural constraints on near-optimal probabilistic inference. (a) Schematic diagram of a random feedforward network, where the input-to-hidden layer weights were set randomly and fixed. (b) Performance of random feedforward networks. The fractional RMSEs in the cue combination task are over 100%. The y-axis is cut off at 100%. (c) Schematic diagram of a recurrent EI network that obeys Dale's law: inhibitory connections are represented by the red arrows and excitatory connections by the blue arrows; inhibitory neurons are represented by lighter colors, excitatory neurons by darker colors. Both input and hidden layers are divided into excitatory-inhibitory sub-populations at a 4-to-1 ratio. Inputs to the network were presented over 10 time steps. (d) Performance of recurrent EI networks.

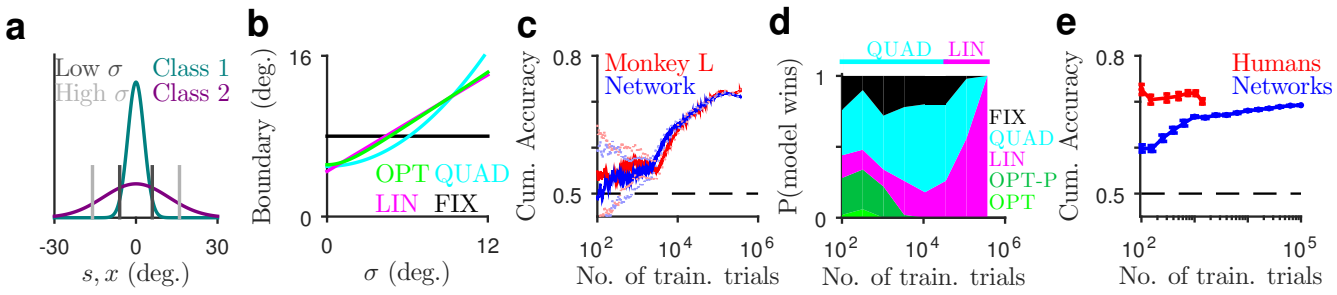

Figure 9: Error-based learning in generic neural networks explains a monkey subject's performance, but not human subjects' performance in a probabilistic binary categorization task involving arbitrary categories. (a) Structure of the two categories. Vertical lines indicate the optimal decision boundaries at low and high noise, σ. (b) Dependence of the decision boundary on sensory noise, σ, for the optimal model (OPT) and three suboptimal heuristic models. (c) Cumulative accuracy of monkey L (red) compared with the cumulative accuracy of a neural network trained with the same set of stimuli (blue). Cumulative accuracy at trial t is defined as the accuracy in all trials up to and including t. The network was trained fully on-line. The input noise in the network was matched to the sensory noise estimated for the monkey and the learning rate was optimized to match the monkey's learning curve (*Methods*). Dotted lines indicate the 95% binomial confidence intervals. (d) The overbar shows the winning models for the monkey's behavioral data throughout the course of the experiment according to the AIC metric. The QUAD model initially provides the best account of the data. The LIN model becomes the best model after a certain point during training. The area plot below shows the fractions of winning models, as measured by their AIC scores, over 50 neural network subjects trained with the same set of input noise and learning rate parameters as the one shown in (c). Similar to the behavioral data, early on in the training, QUAD is the most likely model to win; LIN becomes the most likely model as training progresses. OPT-P is equivalent to an OPT model where the prior probabilities of the categories are allowed to be different from 0.5 (in the experiment, both categories were equally likely). (e) Average performance of 6 human subjects in the main experiment of [38] (red), average performance of 30 neural networks whose input noise was set to the mean sensory noise estimated for the human subjects (blue). Error bars represent standard errors across subjects. Human and monkey behavioral data in (c) and (e) were obtained from [38].

the FIX model assumes that the decision boundary is independent of σ, whereas LIN and QUAD models assume that the decision boundary is a linear or quadratic function of σ, respectively [38]. The learning curve of a monkey subject who performed a large number of trials in this task (monkey L in [38]) is shown in Figure 9 together with the performance of a neural network that received the same sequence of trials as the subject. The input noise of the network was matched to the sensory noise estimated for the subject and the learning rate of the network was optimized to fit the learning curve of the subject. The neural network was trained online, updating its parameters after each trial, in an analogous manner to how the monkey subject learns the task.

Besides providing a good fit to the learning curve of the subject (Figure 9c), the neural networks also correctly predicted the progression of the models that best fit the subject's data, i.e. early on in the training the QUAD model, then the LIN model (Figure 9d). When we performed the same type of analysis on human subjects' data, human observers consistently outperformed the networks and the networks failed to reproduce the learning curves of the subjects (Figure 9e). There might be several possible non-exclusive explanations for this finding. First, prior to the experiment, human observers were told about the task, including what examples from each category looked like. This type of knowledge would be difficult to capture with error-based learning alone and might have given human observers a head-start in the task. Second, human observers might have benefited from possible prior familiarity with similar tasks or stimuli. Third, human observers might be endowed with more powerful computational architectures than simple generic neural networks that allow them to learn faster and generalize better [15].

Low computational complexity of standard psychophysical tasks and the efficiency of generic networks. For each of our tasks, we empirically determined the minimum number of hidden units, n^*, required to achieve a given level of performance (15% information loss for visual search, 10% fractional RMSE or information loss for the other tasks) as a function of the total number of input units, d, in our generic networks. An example is shown in Figure 10a for the causal inference task with $d = 20$ and $d = 220$. The scaling of n^* with d was better

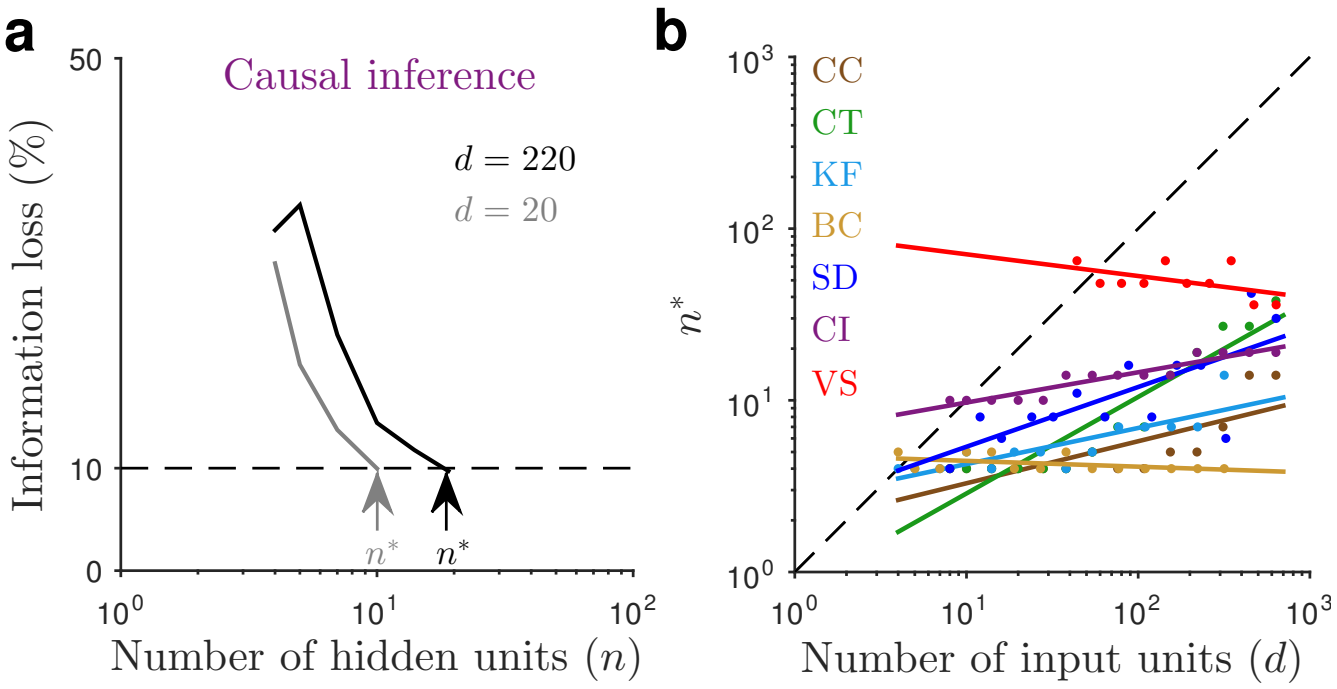

Figure 10: Low computational complexity of standard psychophysical tasks and the efficiency of generic networks. (a) For each number of input units, d, the minimum number of hidden units required to reach within 10% of optimal performance (15% for visual search), n^*, is estimated (shown by the arrows). An example is shown here for the causal inference task with $d = 20$ and $d = 220$ input units, respectively. (b) n^* plotted as a function of the total number of input units, d, in different tasks. Solid lines show linear fits. In these simulations, the number of training trials for each task was set to the maximum number of training trials shown in supplementary Figure S1.

than $O(d)$, i.e. sub-linear, in all our tasks (Figure 10b). Previous theoretical work suggests that this result can be explained by the smoothness properties of the target functions and the efficiency of the generic neural networks with adjustable hidden units. In particular, Barron [1] showed that the optimal number of hidden units in a generic neural network with a single layer of adjustable hidden units scales as $C_f(d)/\sqrt{d}$ with d, where C_f is a measure of the smoothness of the target function, with more smooth functions having lower C_f values. As an example, in [1], it was shown that for the d-dimensional standard Gaussian function, C_f can be upper-bounded by $\sqrt{d}$, leading to an estimate of $O(1)$ hidden units in terms of d. For some of our tasks (e.g. binary categorization; Figure 10b), the scaling of n^* with d was approximately constant over the range of d values tested, suggesting smoothness properties similar to a d-dimensional standard Gaussian. For the other tasks, the scaling was slightly worse, but still sub-linear in every case: in the worst case of coordinate transformation, linear regression of $\log n^*$ on $\log d$ yields a slope of 0.56 ($R^2 = 0.88, P < 10^{-6}$). We can gain some intuition about the relatively benign smoothness properties of our tasks by looking at the analytic expressions for the corresponding target functions (Equations 5-17): although the inputs are high dimensional, the solutions can usually be expressed as smooth functions of a small number of one-dimensional linear projections of the inputs.

The efficiency of our generic networks contrasts sharply with the inefficiency of the manually crafted networks in earlier PPC studies [3, 4, 28–30, 38]: except for the linear cue combination task, these hand-crafted networks used a quadratic expansion which requires at least $O(d^2)$ hidden units. Moreover, unlike generic neural networks, these networks with hand-crafted hidden units are not guaranteed to work well in the general case, if, for example, the target function is not expressible in terms of a quadratic expansion. Stacking the quadratic expansions hierarchically to make the networks more expressive would make the scaling of the number of hidden units with d even worse (e.g. [29]). The fundamental weakness of these hand-crafted networks is the same as that of the random networks reviewed above: they essentially use a fixed basis set theoretically guaranteed to have much worse approximation properties than the adjustable basis of hidden units used in our generic networks [1].

Discussion

We have shown that small generic neural networks trained with a standard error-based learning rule, but without any explicitly probabilistic feedback or training objective, implement probabilistic inference in simple psychophysical tasks and generalize successfully beyond the conditions they are trained in.

For tasks with continuous outputs, we trained our networks to minimize the squared error loss function, which is minimized by the posterior mean estimate. Given the universal approximation guarantees for multilayer neural networks with rectified linear hidden units [26], it is not surprising that our networks can approximate the posterior mean given enough hidden units and training data. However, the findings that near-optimal performance can be achieved even in small networks trained with a relatively small number of training examples and that the networks can generalize successfully beyond the training data they receive depend on the particular problems we studied, in particular, on their low-dimensional nature and their smoothness properties, hence are not predicted by the universal approximation results. Moreover, representing the posterior mean is necessary, but not sufficient for general probabilistic computation: it is also necessary to represent uncertainty on a trial-by-trial basis. Using modular tasks, we showed that the networks implicitly represent uncertainty as well. This finding also holds when the networks are trained with absolute error loss (supplementary Figure S3a).

For tasks with categorical outputs, the output layer of a multilayer neural network is asymptotically guaranteed to converge to the posterior probabilities of the classes under a broad class of loss functions [18], including cross-entropy and mean squared error (supplementary Figure S3b). For the particular problems we studied, our results again show empirically that this convergence can be achieved relatively fast and does not require large networks.

Our tasks all assumed psychophysically realistic levels of sensory noise. At these noise levels, simple heuristic non-probabilistic models that did not take trial-to-trial uncertainty into account are unable to mimic the performance of optimal probabilistic models.

Random networks with only trainable readout weights performed poorly in our tasks. This can be understood as a consequence of the poor approximation properties of such networks [1]. On the other hand, making our networks biologically more plausible by making them fully recurrent and introducing separate excitatory and inhibitory populations throughout the network that respect Dale's law in their connectivity pattern did not significantly impair the performance of the networks.

Relationship to previous work. Our work is consistent with the probabilistic population coding framework according to which, by virtue of the variability of their responses, neural populations encode probability distributions rather than single estimates and computations with probability distributions can be carried out by suitable operations on the corresponding neural responses [28, 48]. However, our work disagrees with the existing literature on the implementation of such computations. We show that these computations do not require any special neural operations or network architectures than the very generic ones that researchers have been using for decades in the neural network community [39, 45, 49].

The recent literature on probabilistic population coding respects the principle that Poisson-like neural variability be preserved between the input and output of a network, because this leads to a fully modular code that can be decoded with the same type of decoder throughout the network [28]. To obtain actual networks, these studies then postulate a literal, one-to-one correspondence between the required neural computations that must be computed at the population level and the computations individual neurons perform. This literal interpretation leads to inefficient neural architectures containing intermediate neurons that are artificially restricted to summing or multiplying the activities of at most two input neurons and that perform substantially different operations in different tasks (e.g. linear summation, multiplication or different forms of divisive normalization) [3, 4, 28–30, 38].

Our generic networks are not necessarily inconsistent with the principle of the preservation of Poisson-like variability between the input and output of a network. Our categorical networks already satisfy this principle, and our continuous networks satisfy it if, instead of a purely linear decoder, we use a linear decoder that is then normalized by the total activity in the hidden layer (supplementary Figure S4). However, our results show that it is unnecessary and inefficient to postulate a direct correspondence between population-level and individual-neuron computations: standard neural networks with rectified linear hidden units that perform the same type of operation independent of the task implement population-level computations required for optimal probabilistic inference far more efficiently.

Experimental predictions. Our results lead to several experimentally testable predictions. First, for gain-

invariant tasks, we predict a novel sparsity-based coding of posterior uncertainty in cortical areas close to the behavioral readout. Stimulus manipulations that increase sensory reliability such as an increase in contrast of the stimulus would be expected to increase the sparseness of the population activity in these areas. Another straightforward consequence of this relationship would be a positive correlation between the performance of the animal in the task and the population sparsity of neurons recorded from the same areas. Second, for gain-dependent tasks such as causal inference, we predict a different coding of posterior uncertainty based on the mean activity in areas close to the readout. Moreover, based on our mean-field model of the mechanism underlying these two types of codes, we expect a trade-off between them: the stronger the correlation between sparsity and posterior uncertainty, the weaker the relationship between the mean activity and posterior uncertainty and vice versa. This can be tested with population recordings from multiple areas in multiple tasks. Third, at the level of single cells, we predict tuning curve sharpening with increased input gain in tasks where a sparsity-based coding of reliability is predicted (Figure 7a). Such tuning curve sharpening has indeed been observed in cortical areas MT [6], MST [20] and MSTd [35]. On the other hand, we expect the input gain to act more like a multiplicative factor in tasks where a mean activity based coding of reliability is predicted (Figure 7b).

Sparse and reliable neural responses have been observed under natural stimulation conditions [9, 17, 44]. Inhibitory currents have been shown to be crucial in generating such sparse and reliable responses [9, 16], reminiscent of the negative mean input requirement in our mean-field model of the sparsity-based coding of posterior uncertainty (Figure 5b).

Limitations. Our networks are highly idealized models of real neural circuits. Although we validated our basic results using biologically more realistic recurrent excitatory-inhibitory networks (Figure 8c-d), even these networks are simplistic models that ignore much of the complexity of real neural circuits. For example, real neural circuits involve several morphologically and physiologically distinct cell types with different connectivity patterns and with potentially distinct functions [19]. Real neurons also implement a diverse set of complicated non-linearities, unlike the simple rectification non-linearity we assumed in our networks. It remains to be determined what possible functional roles this diversity plays in neural circuits.

However, even seemingly drastic simplifications can, in some cases, teach us important insights about the brain. For example, cortical networks are usually highly recurrent, thus modeling them as feed-forward networks might seem like a useless simplification. However, networks with feedback connections can sometimes behave effectively like a feedforward network [14, 36]. As another example, feedforward networks also currently provide the best characterization of the neural responses in higher visual cortical areas [7, 47], even though these areas are known to involve abundant feedback connections both within the same area and between different areas. Therefore, insights gained from understanding simplified models can be quite relevant for understanding real cortical circuits.

Second, our networks were trained with the backpropagation algorithm, which is usually considered to be biologically unrealistic due to its non-locality. Although the backpropagation algorithm in its standard form we have implemented is indeed biologically unrealistic, biologically plausible approximations to backpropagation have been put forward recently [5, 27]. Therefore, it is quite likely that one need not compromise the power of backpropagation in order to attain biologically plausibility.

Third, the stimuli that we used were far from naturalistic. However, the computations required in our tasks capture the essential aspects of the computations that would be required in similar tasks with natural stimuli. Using simple stimuli allows for the parametric characterization of behavior and makes the derivation of the optimal solution more tractable. We have shown here that new insights can be obtained by combining analytically derived optimal solutions with neural networks. For example, understanding the novel sparsity-based representation of posterior uncertainty in the hidden layers of the networks in some tasks but not in others relied on the analysis of the optimal solutions in different tasks.

Finally, as exemplified by the inability of error-based learning to account for the performance of human observers in the binary categorization task, we do not expect error-based learning in generic neural networks to fully account for all aspects of the performance of human observers, and possibly non-human observers as well, even in simple tasks. Relatively mundane manipulations such as changing the target or the distractors, or the number of distractors in a visual search task, changing the duration of a delay interval in a short-term memory task require wholesale re-training of generic neural networks, which seems to be inconsistent with the way human observers, and possibly non-human observers, can effortlessly generalize over such variables. More powerful architectures that combine a neural network controller with an external memory can both learn faster and generalize better [15], offering a promising direction for modeling the generalization patterns of observers in simple psychophysical tasks.

Methods

Data and code availability: The data and code used to obtain the results reported in this paper are available at the following public repository: `https://github.com/eminorhan/inevitable-probability`.

Neural networks: In all networks, the input units were independent Poisson neurons: $\mathbf{r}_{\text{in}} \sim \text{Poisson}(\mathbf{f}(s, c))$ where $\mathbf{f}$ is the vector of mean responses (tuning functions), s is the stimulus and c is a stimulus contrast or coherence variable that controls the quality of sensory information. For the main results presented in the paper, we assume that the effect of c can be described as a multiplicative gain scaling: $\mathbf{f}(s, c) = g(c)\mathbf{f}(s)$ where the individual tuning functions comprising $\mathbf{f}(s)$ were either linear (stimulus demixing), von Mises (visual search) or Gaussian (all other tasks).

To demonstrate the generality of our approach, we also considered two alternative ways in which stimulus contrast or coherence can affect the responses of the input population. In particular, for the cue combination task, we considered tuning functions where c did not have a purely multiplicative effect, but affected the baseline responses as well [11]: $f(s, c) = cf(s) + (1 - c)\beta$ with $0 \leq c \leq 1$ where β was chosen such that the mean response of the input population was independent of c. Secondly, again for the cue combination task with two cues, we considered tuning functions where stimulus contrast c affected both the peak response and the preferred stimuli of input neurons, as reported in [24] for speed tuning in area MT: $f(s, c) = r_0 + Ag(c) \exp\Big(- \frac{1}{2\sigma^2} \big(\log \frac{s+s_0}{Bg(c)\phi+s_0} \big)^2 \Big)$ with the following parameters: $r_0 = 0.5$, $A = 5$, $s_0 = 1$, $\sigma = 1$, $B = 10$ and $g(c) = 1/\big((\alpha c)^{-\beta} + \gamma\big)$ with $\alpha = 10$, $\beta = 2$, $\gamma = 3$. The results for these two cases are shown in Figure 4c.

The hidden units in both feed-forward and recurrent networks were rectified linear units. In feed-forward networks, the hidden unit responses are described by the equation: $\mathbf{r}_{\text{hid}} = [\mathbf{W}_{\text{in}}\mathbf{r}_{\text{in}} + \mathbf{b}]_+$ and in recurrent networks by the equation: $\mathbf{r}_{\text{hid},t+1} = [\mathbf{W}_{\text{in}}\mathbf{r}_{\text{in},t+1} + \mathbf{W}_{\text{rec}}\mathbf{r}_{\text{hid},t} + \mathbf{b}]_+$, where $\mathbf{W}_{\text{in}}$ and $\mathbf{W}_{\text{rec}}$ are the input and recurrent weights respectively and $[\cdot]_+$ denotes elementwise rectification. For tasks with continuous output variables, the network output corresponds to a linear combination of the hidden unit responses: $y = \mathbf{w}^\top \mathbf{r}_{\text{hid}} + b$, and in tasks with categorical variables, the network output was given by a linear combination of the hidden unit responses passed through a sigmoid nonlinearity: $y = \sigma(\mathbf{w}^\top \mathbf{r}_{\text{hid}} + b)$.

In the recurrent EI networks, all connections were constrained to satisfy Dale's law, as described in [41]. The ratio of excitatory to inhibitory neurons in both input and recurrent populations was 4 to 1. Inputs were presented over 10 time steps, but the total input information was equated to the total input information in the feedforward networks.

In random networks, input-to-hidden layer weights were sampled from a normal distribution with zero mean and standard deviation $\sqrt{2/(n + d)}$, where d is the number of input neurons and n is the number of hidden neurons [12]; biases of the hidden units were sampled from a normal distribution with zero mean and standard deviation of 0.1.

In the main simulations, the networks had 200 hidden units. In cue combination, modular cue combination, coordinate transformation, Kalman filtering, binary categorization and causal inference tasks, there were 50 input neurons per input population. To make our results comparable to earlier results, we used 20 input neurons per input population in the visual search task and 10 input neurons per input population in the stimulus demixing task.

Non-probabilistic models: For the Kalman filtering and all cue combination tasks, we used heuristic, non-probabilistic reference models that estimated the individual cues (and the past state in Kalman filtering) optimally, but combined them suboptimally by weighting them equally regardless of their reliability. Note that these models still performed a non-trivial probabilistic computation, namely marginalizing out a nuisance variable, i.e. the input gain, to come up with the optimal estimate of the individual cues. We also note that for the coordinate transformation task, unlike in cue combination or Kalman filtering, the optimal combination rule does not depend on the reliabilities of the individual inputs. For the categorical tasks, we also used non-probabilistic reference models that assumed equal reliabilities for individual inputs. To give the reference models the best chance to perform well, we optimized the assumed common reliability of the inputs by minimizing the information loss.

Training procedure: The feed-forward networks were trained with the standard backpropagation algorithm [39]. The recurrent networks were trained with backpropagation through time [45]. We used the Adam stochastic gradient descent algorithm [22] with learning rate 0.0003 to implement backpropagation. The batch sizes for the updates were 10 for binary categorization, 500 for visual search and 100 for the other tasks. In Figure 9, we used an online

vanilla stochastic gradient descent algorithm with learning rate decrease over trials described by $\eta_0/(1+\gamma t)$ where t is the trial number. The parameters η_0 and γ were fit to the monkey's learning curve.

Training conditions: The "all **g**" conditions in different tasks were as follows. In cue combination and coordinate transformation tasks, all 25 pairs of the form (g_1, g_2) with $g_1, g_2 \in \{0.25, 0.5, 0.75, 1, 1.25\}$ were presented an equal number of times. In Kalman filtering, g was uniformly drawn between 0.3 and 3 at each time step. In binary categorization, the six gain values, $g \in \{0.37, 0.9, 1.81, 2.82, 3.57, 4\}$, were presented an equal number of times. These gain values were calculated from the mean noise parameter values reported for the human subjects in [38]. In causal inference, all 25 pairs of the form (g_1, g_2) with $g_1, g_2 \in \{0.5, 1, 1.5, 2, 2.5\}$ were presented an equal number of times. In stimulus demixing, following [4], c was uniformly and independently drawn between 2 and 9 for each source. In visual search, g was randomly and independently set to either 0.5 or to 3 for each stimulus.

The "restricted **g**" conditions in different tasks were as follows. In cue combination and coordinate transformation tasks, the two pairs $(g_1, g_2) \in \{(0.25, 0.25), (1.25, 1.25)\}$ were presented an equal number of times. In Kalman filtering, g was randomly and independently set to either 0.3 or to 3 at each time step. In binary categorization, g was always 4.2. This gain value corresponds to 100% contrast as calculated from the mean noise parameter values for the human subjects reported in [38]. In causal inference, pairs of the form $(g_1, g_2) \in \{(0.5, 0.5), (2.5, 2.5)\}$ were presented an equal number of times. In stimulus demixing, c was either set to 2 for all sources or else set to 9 for all sources. Similarly, in visual search, g was either set to 0.5 for all stimuli or else set to 3 for all stimuli.

Mean-field model of hidden unit responses: For a given input activity $\mathbf{r}$, we consider the responses of the hidden units as realizations of a random variable r_{hid}. The output weights are also assumed to be realizations of a random variable w. We further assume that w and r_{hid} are independent. The network's output is then proportional to $\langle w \rangle \langle r_{\text{hid}} \rangle$. We want to make this expression invariant to the input gain g. We first introduce a measure of the total sensitivity of this expression to variations in g. We will do this by computing the magnitude of the derivative of $\langle w \rangle \langle r_{\text{hid}} \rangle$ with respect to g and integrating over a range of g values, but we first note that the output weights are already gain invariant, hence we can just consider $\langle r_{\text{hid}} \rangle$. We now have to find an expression for $\langle r_{\text{hid}} \rangle$. The net input to a typical hidden unit is given by:

$$g\mathbf{w}_{\text{in}}^{\top}\mathbf{r} + b \sim \mathcal{N}(\mu_* \equiv g\mu + \mu_b, \sigma_*^2 \equiv g^2\sigma^2 + \sigma_b^2) \tag{1}$$

where $\mathbf{w}_{\text{in}}$ are the input weights to a typical hidden unit. Then:

$$\bar{\mu} \equiv \langle r_{\text{hid}} \rangle = \langle [g\mathbf{w}_{\text{in}}^{\top}\mathbf{r} + b]_+ \rangle = \Big[1 - \Phi\Big(\frac{-\mu_*}{\sigma_*}\Big)\Big]\mu_* + \phi\Big(\frac{-\mu_*}{\sigma_*}\Big)\sigma_* \tag{2}$$

where $\Phi(\cdot)$ and $\phi(\cdot)$ are the cdf and the pdf of the standard Gaussian distribution. As mentioned above, we then introduce the following measure of the total sensitivity of $\bar{\mu}$ to variations in g:

$$T_{\text{var}} = \int_{g_{\min}}^{g_{\max}} |\bar{\mu}'(g)| dg \tag{3}$$

where the prime represents the derivative with respect to g. Because g always appears as $g\mu$ or $g\sigma$ in $\bar{\mu}$ (Equation 2), the parametrization in terms of g, μ and σ is redundant. We therefore set $\sigma = 1$, and hence expressed everything in terms of the scale of σ. We then minimized T_{var} numerically with respect to μ, μ_b and σ_b subject to the constraint that the mean response across different gains be equal to some positive constant K:

$$\frac{1}{g_{\max} - g_{\min}} \int_{g_{\min}}^{g_{\max}} \bar{\mu}(g) dg = K > 0 \tag{4}$$

This ensures that the degenerate solution where the hidden layer is completely silent is avoided.

Task details: In the linear cue combination task, the objective was to combine two cues, $\mathbf{r}_1$ and $\mathbf{r}_2$, encoding information about the same variable, s, in a statistically optimal way. Assuming a squared error loss function, this can be achieved by computing the mean of the posterior $p(s|\mathbf{r}_1, \mathbf{r}_2)$. For a uniform prior distribution, the posterior mean is given by an expression of the form [28]:

$$\hat{s}_{opt} = \frac{\phi^{\top}(\mathbf{r}_1 + \mathbf{r}_2)}{\mathbf{1}^{\top}(\mathbf{r}_1 + \mathbf{r}_2)} \tag{5}$$

where ϕ is the vector of preferred stimuli of input neurons and $\mathbf{1}$ is a vector of ones. This expression is approximate when the prior is not uniform over the entire real line and the quality of the approximation can become particularly bad in the high noise regime considered in this paper. Thus, in practice, we computed posterior means numerically, rather than using the above equation. The equation is still useful, however, in helping us understand the type of computation the network needs to perform to approximate optimal probabilistic inference. During training, the two cues received by the input populations were always non-conflicting: $s_1 = s_2 = s$ and the gains of the input populations varied from trial to trial. The network was trained to minimize the mean squared error between its output and the common s indicated by the two cues.

In the coordinate transformation task, the eye-centered location of an object in 1-d, s_1, was encoded in a population of Poisson neurons with responses $\mathbf{r}_1$ and the current eye position, s_2, was similarly encoded in a population of Poisson neurons with responses $\mathbf{r}_2$. The goal was to compute the head-centered location of the object, which is given by $s = s_1 + s_2$. Assuming uniform priors, the optimal estimate of s can be expressed as [4]:

$$\hat{s}_{\text{opt}} = \frac{\mathbf{r}_1^\top \mathbf{B} \mathbf{r}_2}{\mathbf{r}_1^\top \mathbf{A} \mathbf{r}_2} \tag{6}$$

for suitable matrices $\mathbf{B}$ and $\mathbf{A}$ (see [4] for a full derivation). Again, when the priors are not uniform over the real line, this expression becomes only approximate and posterior means are computed numerically.

In the Kalman filtering task, we considered a one-dimensional time-varying signal evolving according to: $s_t = (1-\gamma)s_{t-1} + \eta_t$, where $\eta_t \sim \mathcal{N}(0, \sigma_\eta^2)$ with $\gamma = 0.1$ and $\sigma_\eta^2 = 1$. At each time t, the stimulus was represented by the noisy responses, $\mathbf{r}_{\text{in},t}$, of a population of input neurons with Poisson variability. The input population projected to a recurrent pool of neurons that have to integrate the momentary sensory information coming from the input population with an estimate of the signal at the previous time step (as well as the uncertainty associated with that estimate) to perform optimal estimation of the signal at the current time step. We decoded the estimate of the signal at each time step by a linear read-out of the recurrent pool: $\hat{s}_t = \mathbf{w}^\top \mathbf{r}_{\text{rec},t} + b$. The network was trained with sequences of length 25 using a squared error loss function. The posterior $p(s_t|\mathbf{r}_{\text{in},1:t})$ is Gaussian with natural parameters given recursively by [4]:

$$\frac{\mu_t}{\sigma_t^2} = \frac{\mu_{\text{in},t}}{\sigma_{\text{in},t}^2} + \frac{(1-\gamma)\mu_{t-1}}{(1-\gamma)^2\sigma_{t-1}^2 + \sigma_\eta^2} \tag{7}$$

$$\frac{1}{\sigma_t^2} = \frac{1}{\sigma_{\text{in},t}^2} + \frac{1}{(1-\gamma)^2\sigma_{t-1}^2 + \sigma_\eta^2} \tag{8}$$

where $\mu_{\text{in},t}$ and $\sigma_{\text{in},t}^2$ are the mean and variance of $p(s_t|\mathbf{r}_{\text{in},t})$ which represents the momentary sensory evidence encoded in the input population. These are, in turn, given by $\mu_{\text{in},t} = \phi^\top \mathbf{r}_{\text{in},t} / \mathbf{1}^\top \mathbf{r}_{\text{in},t}$ and $\sigma_{\text{in},t}^2 = \sigma_f^2 / \mathbf{1}^\top \mathbf{r}_{\text{in},t}$.

In the binary categorization task, the goal was to classify a noisy orientation measurement into one of two overlapping classes that have the same mean but different variances. Given a noisy activity pattern $\mathbf{r}$ over the input population representing the observed orientation, the posterior probabilities of the two classes can be calculated analytically. The log-likelihood ratio of the two categories is given by [38]:

$$d \equiv \log \frac{p(\mathbf{r}|C=1)}{p(\mathbf{r}|C=2)} = \frac{1}{2}\Big(\log \frac{1+\sigma_2^2 \mathbf{a}^\top \mathbf{r}}{1+\sigma_1^2 \mathbf{a}^\top \mathbf{r}} - \frac{(\sigma_2^2 - \sigma_1^2)(\mathbf{e}^\top \mathbf{r})^2}{(1+\sigma_1^2 \mathbf{a}^\top \mathbf{r})(1+\sigma_2^2 \mathbf{a}^\top \mathbf{r})} \Big) \tag{9}$$

where $\mathbf{e} = \phi/\sigma_f^2$ and $\mathbf{a} = \mathbf{1}/\sigma_f^2$. The posterior probability of the first class is then given by a sigmoidal function of d: $p(C=1|\mathbf{r}) = 1/(1+\exp(-d))$.

In the causal inference task, the goal was to infer whether two sensory measurements are caused by a common source or by two separate sources. The log-likelihood ratio of these two hypotheses is given by [30]:

$$d = \frac{z_{11} z_{21}}{z_{12} + z_{22} + J_s} - \frac{1}{2}\Big[\frac{z_{22} z_{11}^2}{(z_{12}+J_s)(z_{12}+z_{22}+J_s)} + \frac{z_{12} z_{21}^2}{(z_{22}+J_s)(z_{12}+z_{22}+J_s)} - \log\Big(1 + \frac{z_{12} z_{22}}{J_s(z_{12}+z_{22}+J_s)}\Big)\Big] \tag{10}$$

where J_s is the precision of the Gaussian stimulus distribution and:

$$\sigma_f^2 z_{11} = \phi_1^\top \mathbf{r}_1 \quad (11)$$

$$\sigma_f^2 z_{12} = \mathbf{1}^\top \mathbf{r}_1 \quad (12)$$

$$\sigma_f^2 z_{21} = \phi_2^\top \mathbf{r}_2 \quad (13)$$

$$\sigma_f^2 z_{22} = \mathbf{1}^\top \mathbf{r}_2 \quad (14)$$

where σ_f^2 is the common variance of the Gaussian tuning functions of the individual input neurons. ϕ_1 and ϕ_2 are the preferred stimuli of the neurons in the first and second populations respectively. For convenience, we assumed $\phi_1 = \phi_2$. The optimal probability of reporting "same cause" is then simply given as: $p(C = 1|\mathbf{r}_1, \mathbf{r}_2) = 1/(1 + \exp(-d))$.

In the stimulus demixing task, the goal was to infer the presence or absence of different signal sources in a mixture of signals with unknown concentrations. As a concrete example, the signals can be thought of as different odors, and the task would then be to infer the presence or absence of different odors in an odor mixture with unknown concentrations [4]. Following [4], we assumed a linear mixing model:

$$o_i = \sum_k w_{ik} c_k s_k \quad (15)$$

where s_k denotes the presence or absence of the k-th odor source, c_k denotes its concentration, o_i is the concentration of the i-th odorant and w_{ik} is the weight of the k-th odor source in the i-th odorant. The task can then be formalized as the computation of the posterior probability of the presence or absence of each odor source, given noisy responses $\mathbf{r} = \{\mathbf{r}_i\}_{i=1}^{n_o}$ of populations of Poisson neurons encoding the odorants: i.e. $p(s_k = 1|\mathbf{r})$. The input populations were assumed to have linear tuning for the odorants: $\mathbf{r}_i \sim \text{Poisson}(o_i \mathbf{f}_i + \mathbf{b}_i)$, where $\mathbf{f}_i$ and $\mathbf{b}_i$ were random vectors with positive entries [4]. As in [4], we assumed four sources and four odorants. The networks were trained to minimize the cross-entropy between the network's outputs and the correct source present/absent labels, s_k.

In the visual search task, the goal was to infer the presence or absence of a target stimulus s_T among a set of heterogeneous distractors. The log-likelihood ratio of the target presence is given by [29]:

$$d = \log \frac{1}{N} \sum_{i=1}^{N} \exp(d_i) \quad (16)$$

where N is the number of stimuli on the display (we assumed $N = 4$) and the local target presence log-likelihoods d_i are given by:

$$d_i = \mathbf{h}_i(s_T)^\top \mathbf{r}_i - \log\left(\frac{1}{\pi} \int_0^\pi \exp(\mathbf{h}_i(s_i)^\top \mathbf{r}_i) ds_i\right) \quad (17)$$

For independent Poisson neurons, the stimulus kernel $\mathbf{h}(s)$ is given by $\mathbf{h}(s) = \log \mathbf{f}(s)$, where we assumed von Mises tuning functions for individual input neurons. The integral in the second term on the right hand side was calculated numerically.

Behavioral data: Human and monkey behavioral data used in Figure 9 were obtained from a previously published study [38]. Human behavioral data reported in Figure 9e are from 6 human subjects (one female) who completed the main experiment in [38]. The behavioral data reported in Figure 9c are from monkey L. Only incomplete behavioral data from another monkey that completed the experiment (monkey A) were available. Because data from all trials are needed to obtain a reliable estimate of the subject's learning curve, data from this monkey were not used in the current paper. For further details about the experimental settings, subjects and model fitting, see [38].

References

[1] Barron AR (1993) Universal approximation bounds for superpositions of a sigmoidal function. IEEE Trans Inf Theory 39(3):930-945.

[2] Battaglia PW, Jacobs RA, Aslin RN (2003) Bayesian integration of visual and auditory signals for spatial localization. JOSA 20(7):1391-1397.

[3] Beck, JM et al. (2008) Probabilistic population codes for Bayesian decision making. Neuron 60:1142-1152.

[4] Beck JM, Latham PE, Pouget A (2011) Marginalization in neural circuits with divisive normalization. J Neurosci 31(43):15310-15319.

[5] Bengio Y, Lee D-H, Bornschein J, Lin Z (2015) Towards biologically plausible deep learning. arXiv:1502.04156.

[6] Britten KH, Shadlen MN, Newsome WT, Movshon JA (1993) Responses of neurons in macaque MT to stochastic motion signals. Vis Neurosci 10:1157-1169.

[7] Cadieu CF, Hong H, Yamins DLK, Pinto N, Ardila D, Solomon EA, Majaj NJ, DiCarlo JJ (2014) Deep neural networks rival the representation of primate IT cortex for core visual object recognition. PLoS Comput Biol 10:e1003963.

[8] Caron SJC, Ruta V, Abbott LF, Axel R (2013) Random convergence of afferent olfactory inputs in the Drosophila mushroom body. Nature 497:113-117.

[9] Crochet S, Poulet JFA, Kremer Y, Petersen CCH (2011) Synaptic mechanisms underlying sparse coding of active touch. Neuron 69:1160-1175.

[10] Ernst MO, Banks MS (2002) Humans integrate visual and haptic information in a statistically optimal fashion. Nature 415:429-433.

[11] Fetsch CR, Pouget A, DeAngelis DC, Angelaki DE (2012) Neural correlates of reliability-based cue weighting during multisensory integration. Nat Neurosci 15:146-154.

[12] Glorot X, Bengio Y (2010) Understanding the difficulty of training deep feedforward neural networks. AISTATS 9:249-256.

[13] Glorot X, Bordes A, Bengio Y (2011) Deep sparse rectifier neural networks. AISTATS 15:315-323.

[14] Goldman MS (2009) Memory without feedback in a neural network. Neuron 61(4):621-634.

[15] Graves A, Wayne G, et al. (2016) Hybrid computing using a neural network with dynamic external memory. Nature 538:471-476.

[16] Haider B, Krause MR, Duque A, Yu Y, Touryan J, Mazer JA, McCormick DA (2010) Synaptic and network mechanisms of sparse and reliable visual cortical activity during nonclassical receptive field stimulation. Neuron 65:107-121.

[17] Haider B, Hausser M, Carandini M (2013) Inhibition dominates sensory responses in the awake cortex. Nature 493:97-100.

[18] Hampshire II JB, Perlmutter BA (1990) Equivalence proofs for multilayer perceptron classifiers and the Bayesian discriminant function. In *Proceedings of the 1990 Connectionist Models Summer School* (Touretzky D, et al. eds.), Morgan Kaufmann.

[19] Harris KD, Shepherd GMG (2015) The neocortical circuit: themes and variations. Nat Neurosci 18:170-181.

[20] Heuer HW, Britten KH (2007) Linear responses to stochastic motion signals in area MST. J Neurophysiol 98:1115-1124.

[21] Hillis JM, Watt SJ, Landy MS, Banks MS (2004) Slant from texture and disparity cues: optimal cue combination. J Vis 4(1).

[22] Kingma D, Ba J (2014) Adam: A method for stochastic optimization. arXiv:1412.6980.

[23] Körding K, Beierholm U, Ma WJ, Quartz S, Tenenbaum JB, Shams L (2007) Causal inference in multisensory perception. PLoS ONE 2(9), e943.

[24] Krekelberg B, van Wezel RJA, Albright TD (2006) Interactions between speed and contrast tuning in the middle temporal area: implications for the neural code for speed. J Neurosci 26(35), 8988-8998.

[25] Kwon O-S, Tadin D, Knill DC (2015) A unifying account of visual motion and position perception. Proc Natl Acad Sci USA 112(26):8142-8147.

[26] Leshno M, Lin VY, Pinkus A, Schocken S (1993) Multilayer feed-forward networks with a non-polynomial activation function can approximate any function. Neural Netw 6:861-867.

[27] Lillicrap TP, Cownden D, Tweed DB, Akerman CJ (2016) Random synaptic feedback weights support error backpropagation for deep learning. Nat Commun 7:13276.

[28] Ma WJ, Beck JM, Latham PE, Pouget A (2006) Bayesian inference with probabilistic population codes. Nat Neurosci 9(11):1432-1438.

[29] Ma WJ, Navalpakkam V, Beck JM, Berg Rv, Pouget A (2011) Behavior and neural basis of near-optimal visual search. Nat Neurosci 14(6):783-790.

[30] Ma WJ, Rahmati M (2013) Towards a neural implementation of causal inference in cue combination. Multisens Res 26:159-176.

[31] Makin JG, Fellows MR, Sabes PN (2013) Learning multisensory integration and coordinate transformation via density estimation. PLoS Comput Biol 9:e1003035.

[32] Maloney LT, Mamassian P (2009) Bayesian decision theory as a model of human visual perception: Testing Bayesian transfer. Vis Neurosci 26:147-155.

[33] Mante V, Sussillo D, Shenoy KV, Newsome WT (2013) Context-dependent computation by recurrent dynamics in prefrontal cortex. Nature 503, 78-84.

[34] Merfeld DM, Zupan L, Peterka RJ (1999) Humans use internal models to estimate gravity and linear acceleration. Nature 398:615-618.

[35] Morgan ML, DeAngelis GC, Angelaki DE (2008) Multisensory integration in macaque visual cortex depends on cue reliability. Neuron 59:662-673.

[36] Murphy BK, Miller KD (2009) Balanced amplification: a new mechanism of selective amplification of neural activity patters. Neuron 61(4):635-648.

[37] Neal RM (1996) Bayesian learning for neural networks. Lecture Notes in Statistics No. 118, New York: Springer-Verlag.

[38] Qamar AT, Cotton RJ, George RG, Beck JM, Prezhdo E, et al. (2013) Trial-to-trial, uncertainty-based adjustment of decision boundaries in visual categorization. Proc Natl Acad Sci USA 110(50):20332-37.

[39] Rumelhart DE, Hinton GE, Williams RJ (1986) Learning representations by back-propagating errors. Nature 323:533-536.

[40] Sclar G, Freeman RD (1982) Orientation selectivity in the cat's striate cortex is invariant with stimulus contrast. Exp Brain Res 46:457-461.

[41] Song HF, Yang GR, Wang X-J (2016) Training excitatory-inhibitory recurrent neural networks for cognitive tasks: A simple and flexible framework. PLoS Comput Biol 12(2): e1004792.

[42] Stettler DD, Axel R (2009) Representations of odor in the piriform cortex. Neuron 63:854-864.

[43] Sussillo D, Churchland MM, Kaufman MT, Shenoy KV (2015) A neural network that finds a naturalistic solution for the prediction of muscle activity. Nat Neurosci 18:1025-1033.

[44] Vinje WE, Gallant JL (2000) Sparse coding and decorrelation in primary visual cortex during natural vision. Science 287:1273-1276.

[45] Williams RJ, Zipser D (1995) Gradient-based learning algorithms for recurrent networks and their computational complexity. In: Chauvin Y, Rumelhart DE (eds.) Back-propagation: Theory, Architectures and Applications. Hillsdale, NJ: Erlbaum.

[46] Wolpert DM, Ghahramani Z, Jordan MI (1995) An internal model for sensorimotor integration. Science 269(5232):1880-1882.

[47] Yamins DLK, Hong H, Cadieu CF, Solomon EA, Seibart D, DiCarlo JJ (2014) Performance-optimized hierarchical models predict neural responses in higher visual cortex. PNAS 111(23):8619-8624.

[48] Zemel R, Dayan P, Pouget A (1998) Probabilistic interpretation of population codes. Neural Comput 10:403-430.

[49] Zipser D, Andersen RA (1988) A back-propagation programmed network that simulates response properties of a subset of posterior parietal neurons. Nature 331:679-684.

Supplementary Figures

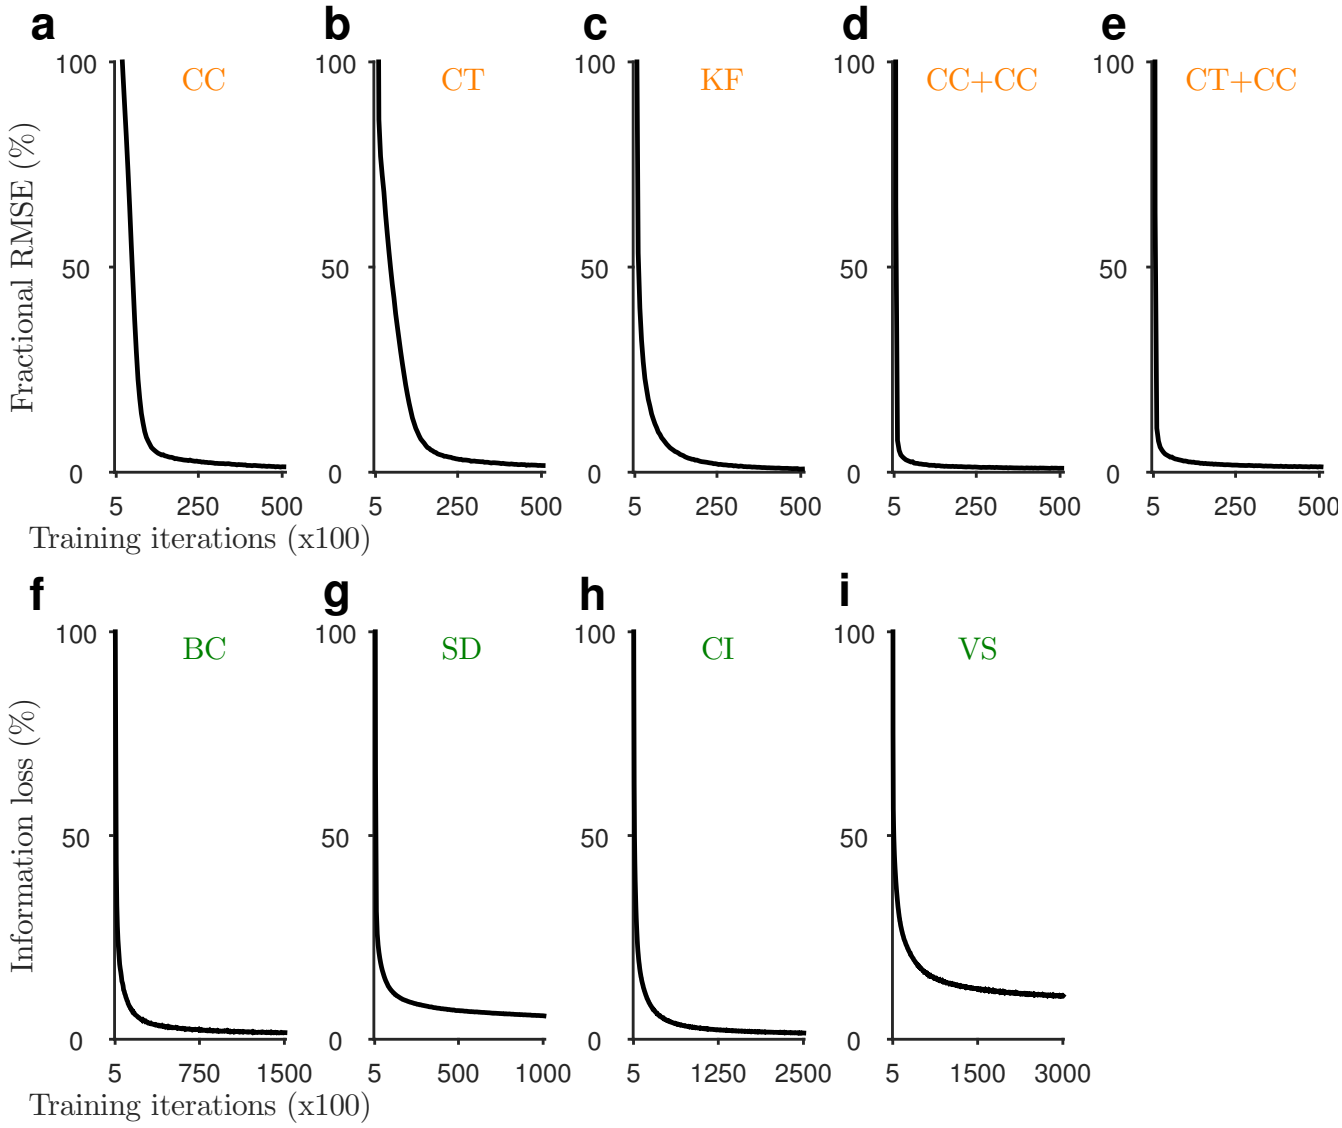

Figure S1: Complete learning curves for all tasks in the "all **g**" conditions. Learning curves are averaged over ten independent runs of the simulations.

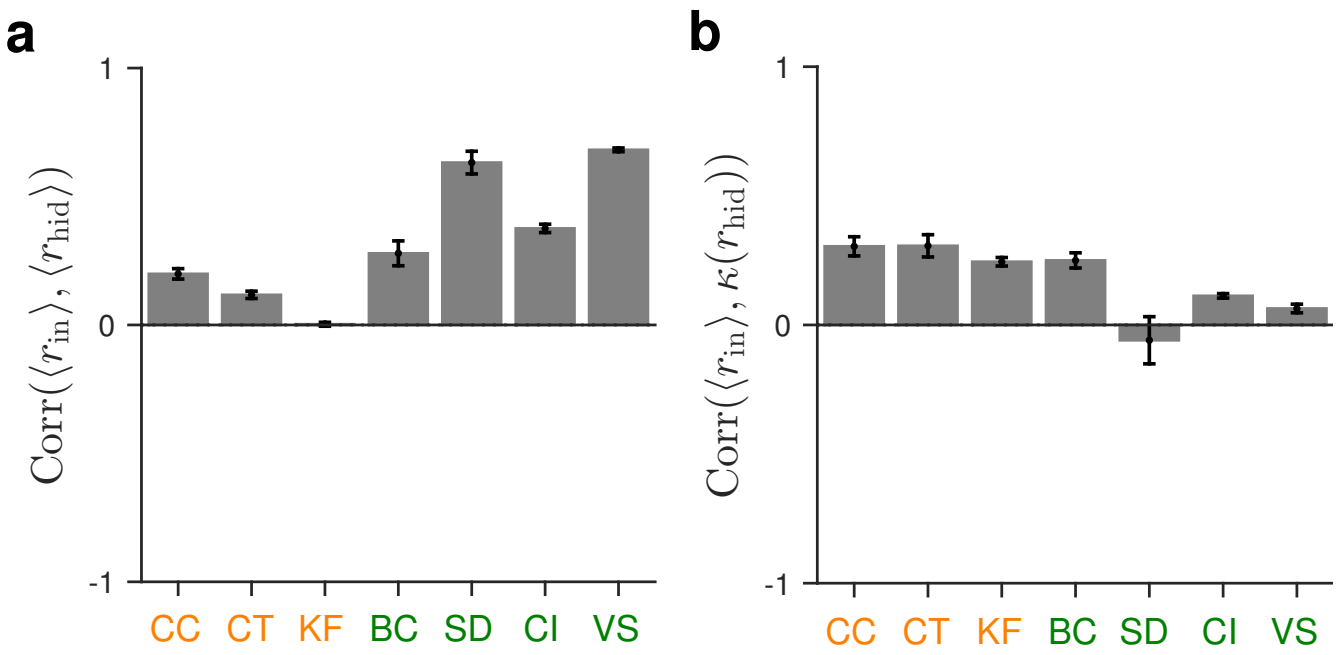

Figure S2: Encoding of posterior uncertainty in the recurrent layer of recurrent EI networks. (a) Correlations between mean input and mean hidden (recurrent) layer responses. (b) Correlations between mean input and sparsity (kurtosis) of hidden layer responses. Mean input was calculated by averaging over both input neurons and time points, whereas mean and kurtosis of hidden layer responses were calculated from the hidden layer responses at the final time point. For the Kalman filtering task, mean input was calculated from the input responses at the penultimate time point and the hidden layer statistics were calculated from the hidden layer responses at the final time point. The pattern of correlations is similar to that observed in fully trained feedforward networks (Figure 6a-b).

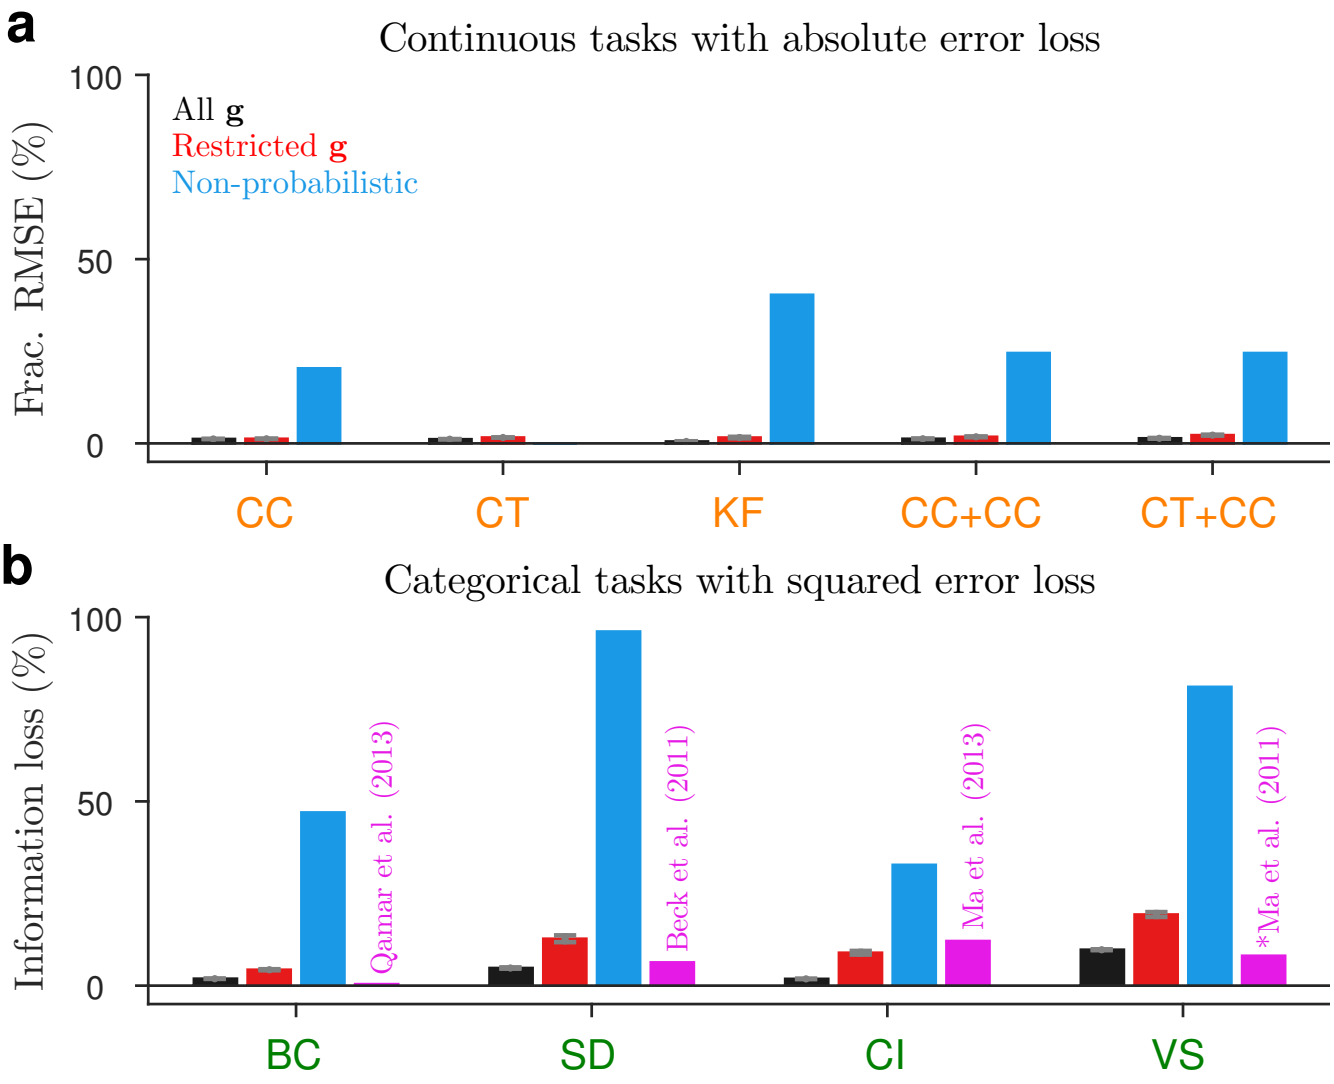

Figure S3: Training with other loss functions: (a) Training in continuous tasks with the absolute error loss. Fractional RMSEs were measured with respect to the posterior median estimates. Note that the networks perform well in the modular tasks, demonstrating the implicit encoding of posterior uncertainty as in mean squared error trained networks. (b) Training in categorical tasks with the squared error loss. As in Figure 2d, magenta bars show the performance of hand-crafted networks reported in earlier works. Performance of well-trained feedforward networks are shown here. Error bars represent standard errors over 10 independent runs of the simulations.

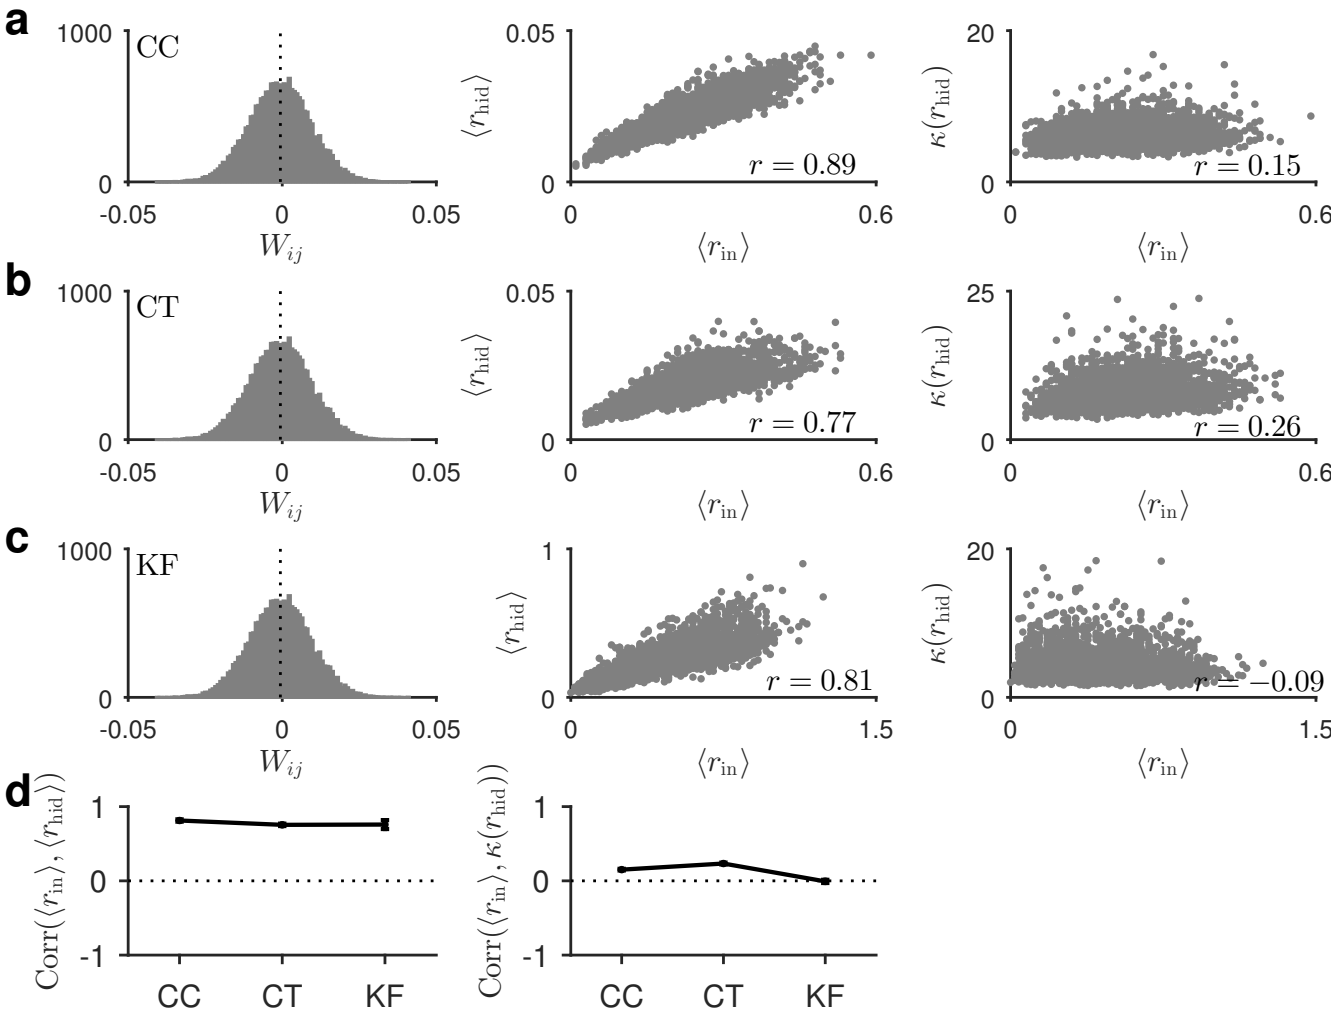

Figure S4: Results for the networks with a divisive normalization decoder (generalized center-of-mass decoder) in the continuous tasks: (a) cue combination, (b) coordinate transformation, (c) Kalman filtering. In these networks, the output of the network was given by $\hat{s} = \mathbf{w}^\top \mathbf{r}_{\text{hid}} / \mathbf{1}^\top \mathbf{r}_{\text{hid}}$. In (a-c), the first column shows the distribution of input-to-hidden layer weights in the trained networks (the distributions are roughly symmetric around zero), the second column is a scatter plot of the mean input vs. the mean hidden layer activity and the third column is a scatter plot of the mean input vs. the kurtosis of hidden layer activity. (d) Correlation between the mean input and the mean hidden unit activity (left) and the correlation between the mean input and the kurtosis of hidden layer activity in the three tasks. Error bars represent standard errors over 10 independent runs of the simulations. In contrast to networks with linear read-outs, networks with a divisive normalization decoder encode posterior uncertainty in the mean hidden layer response, rather than in the sparsity of hidden layer activity.